

Wien, Austria, 7 - 11 September 2009

Proceedings of the 11th European Conference on Computer Supported Cooperative Work

Ellen Balka
Luigina Ciolfi
Carla Simone
Hilda Tellioglu
Ina Wagner

Editors



Springer

ECSCW 2009

ECSCW 2009

Proceedings of the 11th European Conference
on Computer Supported Cooperative Work,
7-11 September 2009, Vienna, Austria

Editors

Ina Wagner

Hilda Telliöglu

Ellen Balka

Carla Simone

Luigina Ciolfi

Editors

Ina Wagner
Vienna University of Technology
Austria

Hilda Tellioğlu
Vienna University of Technology
Austria

Ellen Balka
Simon Fraser University
Canada

Carla Simone
Universita' di Milano-Bicocca
Italy

Luigina Ciolfi
University of Limerick
Ireland

ISBN 978-1-84882-853-7 e-ISBN 978-1-84882-854-4
DOI 10.1007/978-1-84882-854-4
Springer Dordrecht Heidelberg London New York

Library of Congress Control Number: 2009934004

© Springer-Verlag London Limited 2009

Apart from any fair dealing for the purposes of research or private study, or criticism or review, as permitted under the Copyright, Designs and Patents Act 1988, this publication may only be reproduced, stored or transmitted, in any form or by any means, with the prior permission in writing of the publishers, or in the case of reprographic reproduction in accordance with the terms of licenses issued by the Copyright Licensing Agency. Enquiries concerning reproduction outside those terms should be sent to the publishers.

The use of registered names, trademarks, etc., in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant laws and regulations and therefore free for general use.

The publisher makes no representation, express or implied, with regard to the accuracy of the information contained in this book and cannot accept any legal responsibility or liability for any errors or omissions that may be made.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Table of Contents

From the editors	vii
ECSCW 2009 Conference Committee	ix
ECSCW 2009 Program Committee	x
The boundaries of participatory citizenship <i>Nikolaj Gandrup Borchorst, Susanne Bødker and Pär-Ola Zander</i>	1
Research Project as Boundary Object: negotiating the conceptual design of a tool for International Development <i>Ann Light and Theresa Dirndorfer Anderson</i>	21
Supporting Nurses' Information Flow by Integrating Paper and Digital Charting <i>Charlotte Tang and Sheelagh Carpendale</i>	43
The eDiary: Bridging home and hospital through healthcare technology <i>Rikke Aarhus, Stinne Aaløkke Ballegaard and Thomas Riisgaard Hansen</i>	63
PRODOC: an Electronic Patient Record to Foster Process-Oriented Practices <i>Federico Cabitza, Carla Simone and Giovanni Zorzato</i>	85
Finding the other 5%: Understanding the role of social networking technologies in building personal networks for young adults with cancer <i>Shawn Ashkanasy, Frank Vetere, Graeme Shanks and Hilary Davis</i>	105
Trust and Social Capital: Revisiting an Offshoring Failure Story of a Small German Software Company <i>Alexander Boden, Bernhard Nett and Volker Wulf</i>	123
Return On Contribution (ROC): A Metric for Enterprise Social Software <i>Michael Muller, Jill Freyne, Casey Dugan, David Millen and Jennifer Thom-Santelli</i>	143
Collaborative Practices that Support Creativity in Design <i>Dhaval Vyas, Dirk Heylen, Anton Nijholt and Gerrit van der Veer</i>	151
'Talking about (my) Generation?': Creativity, Practice, Technology & Talk <i>David Martin, Jacki O'Neill and Dave Randall</i>	171
Using Annotations in a Collective and Face-to-Face Design Situation <i>Sylvie Guibert, Françoise Darses and JF Boujut</i>	191

Table of Contents

We can work it out: Collaborative Conflict Resolution in Model Versioning <i>Petra Brosch, Philip Langer, Martina Seidl, Konrad Wieland and Manuel Wimmer</i>	207
On the effects of Refactoring in the Coordination of Software Development Activities <i>Cleidson de Souza, Maryanne Rosa, Crys Goto, Jean Costa and Pedro Trecanno</i>	215
Divided by a common acronym: On the fragmentation of CSCW <i>Kjeld Schmidt</i>	223
Collaboration in Metagenomics: Sequence Databases and the Organization of Scientific Work <i>Matthew J. Bietz and Charlotte P. Lee</i>	243
Towards <i>LivingAgendas</i> - Shaping the next generation of business meetings <i>Till Schümmer, Hilda Tellioglu and Jörg M. Haake</i>	263
Analyzing Multimodal Communication around a Shared Tabletop Display <i>Anne Marie Piper and James D. Hollan</i>	283
Status on Display: a Field Trial of Nomadic*Viz <i>Xianghua Ding and Donald Patterson</i>	303
The ‘out-of-avatar experience’: object-focused collaboration in Second Life <i>Greg Wadley and Nicolas Ducheneaut</i>	323
Character Sharing in World of Warcraft <i>Nelson Wong, Anthony Tang, Ian Livingston, Carl Gutwin and Regan Mandryk</i>	343
High-Octane Work: The oil and gas workplace <i>Clint Heyer</i>	363
Asking for the moon: Or model-based coordination in distributed design <i>Kjeld Schmidt, Hilda Tellioglu and Ina Wagner</i>	383
Information Curators in an Enterprise File-Sharing Service <i>Michael Muller, David Millen and Jonathan Feinberg</i>	403
Index of authors	413

From the Editors

This volume represents the proceedings of ECSCW'09, the 11th European Conference on Computer Supported Cooperative Work, held in Vienna, Austria. Each conference offers an occasion to critically review our research field, which has been multidisciplinary and committed to high scientific standards, both theoretical and methodological, from its beginning. Ongoing discussion has identified several challenges, which also become visible in the conference programme. One challenge comes from emerging new technologies connected to 'social computing', gaming, as well as applications supporting citizen participation in their communities. To examine user experiences and collaborative aspects of these applications attracts the interest of many colleagues and also some newcomers to the field and there are some fine studies represented in this conference volume. As boundaries between home and work erode with the increased movement of work into home environments, and new applications further blur the once separate conceptions of work and leisure, our intellectual community faces challenges in the ways we think about and study work. Other challenges result from transformations of the world of work itself and the role of IT in these. They have been taken up in in-depth studies of design practice, software development, and manufacturing, as well as in the growing body of research on health care contexts and applications. In times of rapid societal change and crisis there is a need for examining not only the social relevance of CSCW research topics but also to look into the theoretical and methodological framework, on which this research is based and to try achieve greater conceptual clarity and methodological validity. Finally, there is the question of what is the European perspective in our community and whether it is worthwhile to anchor our research more firmly in such a perspective. Of high relevance to our field is the strong grounding of technology development in an understanding of human activity. In Europe we have a strong philosophical, sociological and anthropological research tradition, on which our community can build when augmenting human practice with new artefacts, media and infrastructures.

The nineteen full papers, four short papers and one discussion paper selected for this conference deal with and reflect on some of these challenges. They form the core of a single-track conference programme which is somehow a tradition in ECSCW. We are also excited about the 10 workshops and masterclasses that cover a broad range of topics and allow for wider and more active participation and will be published in the on-line supplementary proceedings, as well as the demonstrations, videos, and posters.

Many people have worked hard to ensure the success of this conference, and we briefly acknowledge them here: all the authors who submitted high quality papers; all those who contributed through taking part in workshops, masterclasses, demonstrations, and posters; the Programme Committee, which dedicated time and energy to reviewing and discussing individual contributions and shaping the programme; the student volunteers who provided support throughout the event; and all the sponsors and those who offered their support to the conference.

Ina Wagner, Hilda Tellioglu
Ellen Balka, Carla Simone, Luigina Ciolfi

ECSCW 2009 Conference Committee

Conference Chair

Ina Wagner, Vienna University of Technology, Austria

Conference Vice-Chair

Hilda Tellioglu, Vienna University of Technology, Austria

Program Co-Chairs

Ellen Balka, Simon Fraser University, Canada

Carla Simone, Universita' di Milano-Bicocca, Italy

Workshops and Masterclasses Co-Chairs

Giulio Jacucci, Helsinki Institute for Information Technology, Finland

Dave Randall, Manchester Metropolitan University, UK

Demos and Videos Co-Chairs

Alessandra Agostini, Universita' di Milano-Bicocca, Italy

Valérie Maquil, Vienna University of Technology, Austria

Posters Co-Chairs

Pernille Bjørn, ITU Copenhagen, Denmark

Carl Gutwin, University of Saskatchewan, Canada

Doctoral Colloquium Co-Chairs

Giorgio De Michelis, Universita' di Milano-Bicocca, Italy

Kjeld Schmidt, Copenhagen Business School, Denmark

Proceedings Chair

Luigina Ciolfi, University of Limerick, Ireland

Student Volunteers Coordinator

Markus Rester, Vienna University of Technology, Austria

ECSCW 2009 Program Committee

- Mark Ackerman, University of Michigan, USA
Alessandra Agostini, Università di Milano-Bicocca, Italy
Liam Bannon, University of Limerick, Ireland
Pernille Bjørn, ITU Copenhagen, Denmark
Jeanette Blomberg, IBM Almaden Research Center, USA
Claus Bossen, University of Aarhus, Denmark
Geof Bowker, Santa Clara University, USA
Susanne Bødker, University of Aarhus, Denmark
Tone Bratteteig, University of Oslo, Norway
Monika Büscher, Lancaster University, UK
Federico Cabitza, Università di Milano-Bicocca, Italy
John M. Carroll, Pennsylvania State University, USA
Luigina Ciolfi, University of Limerick, Ireland
Françoise Darses, Conservatoire National des Arts et Métiers, France
Giorgio De Michelis, Università di Milano-Bicocca, Italy
Prasun Dewan, University of North Carolina, USA
Monica Divitini, Norwegian University of Science and Technology, Norway
Geraldine Fitzpatrick, University of Sussex, UK
Antonietta Grasso, Xerox Research Centre Europe Grenoble, France
Tom Gross, Bauhaus-University Weimar, Germany
Carl Gutwin, University of Saskatchewan, Canada
Christine Halverson, IBM T.J. Watson Research Center, USA
Richard Harper, Microsoft Research Centre Cambridge, UK
Bo Helgesson, Blekinge Institute of Technology, Sweden
Thomas Herrmann, University of Bochum, Germany
Jon Hindmarsh, King's College London, UK
Eva Hornecker, University of Strathclyde, UK
Kori Inkpen, Microsoft Research, USA
Giulio Jacucci, Helsinki Institute for Information Technology, Finland
Marina Jirotko, Oxford University Computing Laboratory, UK
Aditya Johri, Virginia Tech, USA
Kari Kuutti, University of Oulu, Finland
Wendy Kellogg, IBM T.J. Watson Research Center, USA
Timothy Koschmann, Southern Illinois University, USA
Steinar Kristoffersen, Østfold University College, Norway
Charlotte Lee, University of California Irvine, USA
Paul Luff, King's College London, UK
Wayne Lutters, University of Maryland, USA
Gloria Mark, University of California, Irvine USA
David W. McDonald, University of Washington, USA
David Millen, IBM T.J. Watson Research Center, USA
Michael Müller, IBM T.J. Watson Research Center, USA
Keiichi Nakata, University of Reading, UK
Bonnie Nardi, University of California Irvine, USA
Maria Normark, Södertörn University College, Sweden
Jacki O'Neill, Xerox Research Centre Europe Grenoble, France
Volkmar Pipek, University of Siegen, Germany
Wolfgang Prinz, Fraunhofer FIT, Germany
Rob Procter, University of Manchester, UK
Dave Randall, Manchester Metropolitan University, UK
Toni Robertson, University of Technology Sydney, Australia
Yvonne Rogers, Indiana University, USA
Mark Rouncefield, Lancaster University, UK
Pascal Salembier, Université de Technologie de Troyes, France
Kjeld Schmidt, Copenhagen Business School, Denmark
Nadine Schuurman, Simon Fraser University, Canada
Stacy Scott, University of Waterloo, Canada
Dan Shapiro, Lancaster University, UK
Lucy Suchman, Lancaster University, UK
Yngve Sundblad, Royal Institute of Technology, Sweden
Hilda Tellioglu, Vienna University of Technology, Austria
Michael Twidale, University of Illinois Urbana-Champaign, USA
Volker Wulf, University of Siegen, Germany

The boundaries of participatory citizenship

Nikolaj Gandrup Borchorst, Susanne Bødker, Pär-Ola Zander
Department of Computer Science, Aarhus University
(*borchorst, bodker, zander*)@cs.au.dk

Abstract. This paper explores the space between municipal administrative systems and citizens' web use. It addresses the possibilities of drawing new boundaries between public administration and citizens' everyday lives through a shared planning and visualization artifact, embedded into Web 2.0. The case deals with planning, advising and control of parental leave. This process involves several citizens, the municipal office, employers, as well as the laws regulating parental leave, and the collective agreements supplementing this legislation. The municipal office controls that citizens and employers comply with the law. At the same time it is often the only reliable source of overview of the law, and of leave days recorded. This paper analyses the current situation, presents an exploratory design process and outcome, probing the boundaries between citizens and the municipal office. Focusing on boundaries and tribes, the paper discusses how new forms of web technologies may improve communication between citizen and government and facilitate collaborative user empowerment: Participatory citizenship. Where Web 2.0 technology is often thought of as tearing down boundaries between individuals, this case points to the importance of a focus beyond individual users, and a renegotiation of boundaries between citizens and caseworkers in the context of other groups of actors.

Introduction

This paper presents an action research project addressing the practices and technologies of applying for and being counseled about parental leave funding. The setting is a major Danish municipality where expecting parents turn to the municipal office for approval and advice on their parental leave funding scheme. The overall purpose of the project was to explore Web 2.0 technologies for improving cooperation and communication in this field.

In order to understand the context of this study, we take a look at the history and state of Danish municipal government services. Traditionally, the formal interaction between citizens and government has been thought of in terms of one citizen to one service. In 2003 a new citizen to this municipality had to visit 3-4 different physical municipal offices. These visits were required to receive the services needed to become a citizen within the municipal system, and thus qualify for a salary or social benefits (Bødker, 2004). The citizen, in other words, created the connection between municipal services.

For many years the Danish society has been assigning unique identifiers to every citizen at birth. Legislation strongly restrains governmental institutions' cross-use of information, and in practice the IT systems are separate. Accordingly, there is no single point in the public administration where somebody or a certain instance "knows everything about me" (Luitjens, 2008).

The trend in government organization (Wimmer, 2002) has for a few years gone towards "personalized service," and "one-stop-shopping". The focal point of these tendencies has been the idea of the citizen gaining easier access to the whole spectrum of governmental services from one physical location or through one website. Borger.dk is a Danish example of such a website. Such a "personalized service" is challenged when services pertain, not to one, but to several citizens. An example of this is parental leave. Historically the caring of a newborn baby was expected to be a woman's job. Today, however, this job is to a larger and larger extent shared between the parents. Consequently, today, the municipality needs to handle parental leave as a service, which involves several citizens—the mother, the father and the child. These primary stakeholders are surrounded by a web of additional stakeholders consisting of employers, other municipal offices, unions and regulations, relatives and friends of the expecting parents, and not least the circles of (expecting) mothers, organized by the visiting nurse into what is called Mothers' groups. These groups proved to be an important source of information among many parents, and as such they were of interest to our design process.

In the following, we further explore the practices and technologies applied by parents and municipal caseworkers, together or separately, to deal with the issues of planning and controlling the parental leave funding.

Parental leave planning—analysis and design

The parental leave case is part of the eGov+ project, which explores e-governance services and infrastructure. The pivotal idea of the project is to examine how citizens may be supported in achieving as much as possible on their own and in cooperation with other citizens, and how collaboration between citizens and municipal services may be enhanced.

In 2008 we carried out a participatory design-oriented study of practices and technologies of the planning, advising and controlling of parental leave funding, presented here. Over six months we interviewed expecting and new parents. We conducted field studies in the municipal office, including flow-oriented analysis workshops. We carried out iterative, participatory design, prototyping alternatives for shared planning and overview tools. The activities were conducted in the same time order as they are enumerated, with the exception of A4 and A5, which were conducted concurrently (Figure 1).

We recorded empirical data in the form of audio, video and pictures, supplemented with thorough notes.

Activity	Approach	Participants	Prototype
A1	2 days of field study, participant observation	Employees at two municipal offices	
A2	3 workshops on work and document flow	Employees at two municipal offices	
A3	6 hours of workshops	Mothers' groups	
A4	3 sessions of workshops	5 caseworkers	Early paper prototype
A5	7 hours of pluralistic walkthroughs (Bias, 1994)	Citizens (individual mothers)	Early paper prototype
A6	2 hours prototype workshop	Case workers from all city offices	Several alternative prototypes
A7	2 hours prototype workshop	Project members from 3 municipalities	Several alternative prototypes

Figure 1 Overview of activities (activity numbers used as references to quotes in this paper).

As figure 1 indicates, caseworkers participated actively in the design process, as they and their municipal organizations were partners in the process. The caseworkers were to learn from the process and the approaches, which were classical participatory design methods (Greenbaum & Kyng, 1991). As such, our engagement with citizens was more ad-hoc. However, it was in our interest, as well as in the interest of the municipal office, to enroll citizens as volunteers (Gaver et al., 2004), or informants (Scaife & Rogers, 1998).

We have in particular applied three approaches in the project: Participant observation and interviews were used to get an early understanding of the field. The participants in these workshops were users, designers and researchers who together carried out structured activities in order to explore and discuss specific issues such as document flows. Moreover the participants tried out various prototypical solutions in a structured form. In this particular case workshops included focus group interviews with Mothers' groups, walkthroughs of paper prototypes with parents and caseworkers, situation game-inspired discussions of social network and adaptive document technologies with caseworkers, and finally scenario-based, hands-on use of running prototypes with caseworkers (Greenbaum & Kyng, 1991). Iterative prototyping served the purpose of giving future users access to hands-on experience, while making it possible to work with alternatives as opposed to just *one* solution. When applied in a research setting, prototypes may be used to capture and probe research hypotheses while providing

a practical point of departure for discussions among future users (Bødker, 1999). In the eGov+ project, rough and meticulous paper prototypes, as well as running web-based prototypes were used to illustrate our proposed conceptual solution—a timeline.

Throughout the process we have focused on contradictions between the needs of the stakeholders, which manifested themselves as breakdowns in the parental leave practice. For instance we detected a number of situations where parents gave up in their planning process and settled for a simple, but less satisfying alternative. We also probed for (and identified) contradictory drives in organizational goals (in particular correct case processing versus better service).

Outline and idea of the prototypes

The design case deals with planning and advising as well as control of parental leave. As previously mentioned, it involves several stakeholders. The municipal office supervises that citizens and employers comply with the law. At the same time it is often the only reliable source of information regarding the interpretation of the law, as well as the status overview of the number of leave days spent by the parents. This overview is particularly problematic, since parental leave can be taken in a number of ways over a 9-year time period.

There are a limited number of online information sources available to parents, often provided by unions, employers, or in the form of private online communities (e.g., www.navlestrengen.dk). Such privately hosted online communities are mostly regular public fora, not only concerning the issue of parental leave. As such there exists no one place that gives either an overview of the complex legislation and myriad of agreements, or professional guidance to specific cases. The only way to obtain such guidance is by contacting the municipality, the union, or the employer. In addition, the compensation for parental leave comes from several agencies. If the parent is subject to a collective agreement between a labor union and the employer, both the municipality and the employer contribute to the compensation paid. Employers give different compensations, and therefore it is impossible for the municipality to provide accurate information in this area. Essentially, this complex constellation of stakeholders, rules and regulations makes it extremely difficult for both citizens and caseworkers to maintain an overview of the choices to be made by the citizens, and the rules restricting these choices. As a consequence many citizens find themselves frustrated with both legislation and the service provided by their municipality. Our interviews point towards the fact that citizens are essentially fond of the flexibility, which the legislation grants them, while at the same time being fundamentally unhappy with the complexity that this flexibility entails. One citizen explained her frustration with not being able to find a clear answer to her questions:

Quote 1: So we think its been really, really, really hard to figure out the rules. [...] It hasn't been easy [...]. But it's probably because we found ourselves caught in the middle of those collective bargainings for both his and my profession. [Mother, A3]

The fundamental idea of our the design concept was to explore how citizens could help themselves and each other in understanding, planning and applying for parental leave funding. At the same time we wished to enhance the communication between the citizens and the municipal caseworkers, when such communication was needed.

The overall ideas included support to shape and visualize the leave for both parents using a timeline; a shared object of negotiation between citizens and caseworkers; aid in evaluating alternative what-if scenarios in terms of time and money; possibilities of sharing with friends and adding information e.g. from unions; streamlining of the application process by eliminating unnecessary parts of forms and redundant information, some of which would come from other sources such as the employers.

Visualizing the parental leave and the regulations and administrative procedures surrounding it, as a timeline has several advantages. It can function both as a planning tool, showing what is to come, and as a historical overview, showing what has already occurred. As remaining available parental leave is determined by time already spent, the display of previous history is crucial. The timeline cannot be used in isolation with the current regulations. There are still bureaucratic procedures to follow, and actual applications to fill out with a particular timing (Bohøj & Bouvin, 2009). Keeping track of such timely procedures is the second dimension of the timeline (see Figure 1).

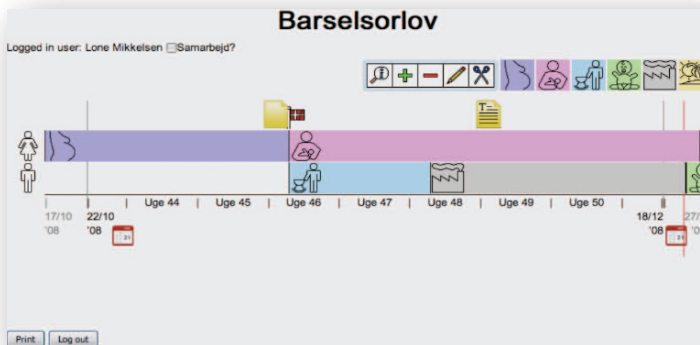


Figure 1 The timeline

The timeline is divided into three horizontal sections. Per default the middle section shows a six-week period centered on the current date. At both ends compact sections show events outside this main section. Users may zoom in or out respectively to render more details visible or to obtain an overview of a longer period. Furthermore, citizens can share their timeline with caseworkers.

Scenario 1 (below) illustrates how the timeline is to be used collaboratively.

It is well established that a timeline can be an appropriate concept for visualizing personal cases in government activity (Plaisant et al., 1996). Our timeline concept is based on it being web-based and sharable between parents, as well as with municipal offices when access is explicitly granted. It flexibly embraces several children and several parents. It makes extraction from municipal systems and generation of necessary documents easy, while still under control of the parent. Furthermore, the timeline can be summed up and shared with friends, or tailored, e.g. by unions or major employers to match information about parental leave conditions pertaining to specific agreements. However, there are many challenges to actually making such complex sharing and collaboration possible. The situation is much more complicated than what is generally believed to be the case within current e-governance: That caseworkers and citizens simply can and should use one and the same web-based solution (KL, 2006, p. 8).

Mette and Jacob are together and have a boy named Magnus, age 5 months.

When Mette and Jacob were first planning their leave, they visited the municipal website to use timeline planning tool. They found a plug-in provided by Jacob's union, which helped them understand his specific situation regarding salary.

Mette has a friend from the University, Anne, who was recently on leave, and Mette asked to have a look at Anne's timeline. Anne shared this with her, and pointed to the fact that several other friends had uploaded anonymized versions of their timelines to a Facebook group, which could be found e.g. via the municipal website.

Once Mette and Jacob have decided on their plan, Mette shares her part with the municipality Jacob places a request for his employer to fill in the necessary information, sign the plan and send it to the municipality.

One day Mette gets a call from an old friend from school. He offers her a job starting in just one month? Mette finds the job very attractive, but the date is earlier than she had planned. Mette and Jacob try to figure out their options: Might Jacob be able to start his leave earlier? How will this affect their budget? Will his employer agree? Maybe he can use some of his vacation instead? They look at the timeline again. Obviously, some things can no longer be changed as Mette has spent 5 months of her leave already. Sitting at the computer they try out various scenarios. They look at what other people have done, by browsing the Facebook group. The sandbox allowing for what-if explorations gives the couple a very good feel for what the legislation allows, what is most beneficial with respect to Jacob's salary agreement, etc.

Once they have decided, Jacob shares the plan with his employer to get approval of the change of plans. So does Mette with the municipality, before accepting the job offer.

Scenario 1 The future use of the timeline

Theoretical framing: Boundaries and tribes

In seeking inspiration from CSCW to address our findings and design, it seems evident that the empirical situation is concerned with the meeting of communities of practices or cultures. Thus future technology ought to support these webs of

actors as opposed to focusing on “within communities” issues. Web 2.0 technologies (O’Reilly, 2005) might in the outset seek to tear down the boundaries between citizens and caseworkers, e.g. by allowing citizens to share electronic web forms from home with the caseworker in the office. Nevertheless this is a much too simplified perspective for a variety of reasons, which we will discuss. Boundary objects (Star & Griesemer, 1989; Star, 1989) are often seen as addressing objects and information that cross boundaries between communities. As discussed by e.g. Lee (2005), this seems oversimplified, and a “sailing across” boundaries is not all that happens. Lee suggests a focus on boundary negotiating artifacts that push or strengthen the boundaries. We have chosen the title of this paper in an attempt to address what it means for citizens and caseworkers to take part, when web technology is placed on the boundaries to municipal services, rather than when web technology aims at tearing down such boundaries.

Bødker et al. (2003) similarly discuss technology for boundaries, and build on Barth’s conception of boundaries in the study of an organizational setting. The main argument of Barth’s perspective is to focus on contexts and situations in which boundaries are generated. The focal point of his claim is that it is the boundaries that define the group rather than the cultural core. Organizational boundaries become visible in organizational structures and rules, and they exist as invisible patterns between individuals and different groups of people. Boundaries, for example, separate one work domain from another, and one profession from another. Or they can be drawn between groups of people defined by shared interests in for example new technology. Boundaries outline the identity of the community and are marked because communities interact with entities from which they are, or wish to be, distinguished (Barth, 1969). The manner in which they are marked depends upon the specific community–administrative boundaries may be statute, cultural boundaries by language.

Bødker et al. (2003) look at computer technology and the way it supports or prohibits boundary work. The authors show how boundaries between the organization and its customers lead to new needs for the sharing of information and computer applications. Within the organization, boundaries between competencies and areas of responsibility are equally influential on the need for information sharing and computer support. Compared with other studies focusing on borders or boundaries, this perspective is more dynamic in that it does not take borders for granted, but argues for technology designed to support the borders, which might exist. Clement & Wagner (1995) similarly look at boundaries that are enforced by organizations, e.g. with the purpose of reducing complexity. They call this fragmented exchange. Neither of these two ways of looking at boundaries primarily focuses on boundary crossing capacities. Rather, they study how technologies may enforce or move boundaries for various reasons, in the same way as Lee’s (2005) boundary negotiating artifacts. In this paper we look at boundaries, the way boundaries get maintained and changed, the roles of technological artifacts in

these boundary negotiations, and in particular the way future technological artifacts may change the landscape.

Taking Barth's lesson that communities are constituted by their boundaries, and that these boundaries are created in the meeting between communities, we look for where communities meet. One such important place is in the objects produced for others to use. However, it should not be taken for granted that these boundaries are neither rationally, nor consciously drawn. Maffesoli's concept of the *tribe* is used to describe how today's society consists of a bricolage of tribes. These are organized by a local common aesthetical experience and localized ethics and customs, such as brand communities, punks, regular commuters, or the village neighbourhood (Maffesoli, 1996). Contrary to the standpoint that the citizen *as an individual* is independently and rationally consuming goods, citizens participate in the consumption of services in a way that cannot be reduced to explanation by an individual, rational purpose in the tribal community. Citizens adjust their interaction with the government in order to satisfy their needs, but their behaviour is seldom strictly rational (March & Simon, 1958). Neither do citizens necessarily act as rational collectives with common purposes (Shield in Maffesoli, 1996, p. xi). In this paper, we use the concept of tribes to differentiate certain citizen communities from more purpose-driven communities. The three main boundary constituents of a tribe, available for empirical inquiry, are aesthetics, ethics and customs.

The *aesthetics* of a community is constituted by collective emotions. Collective emotions are created and shaped through observable interaction; it is not that an individual feels something, which is later externalized. Community rules are what is experienced as "should" or "ought to" by the members of the community: The *ethics*. These ethics also mark the boundaries to those outside the community. Ethics are experienced rules for behaviour, whereas aesthetics are of a more interpretative nature. *Customs* are "*The collection of common usages that allow a social entity to recognize itself for what it is.*" (Maffesoli, 1996, p. 21). Customs are different from ethics in that they are the actual acts, not the experienced rules according to which one ought to act.

We present a setting into which a number of prototypes were introduced to explore and support the communication between a particular kind of citizens (expecting and new parents) and between these citizens and municipal caseworkers. With the above perspective it makes sense to see these prototypes as attempts to explore and develop the boundaries, similarly to the Lee's (2005) description.

The boundaries and tribes of parental leave

In the analysis we focus on the expecting parents on the one hand, and the caseworkers on the other, with a view to the context of other actors and artifacts such as the legislation, agreements and municipal IT-systems. With an understanding

of these groups as *tribes*, we look at the boundaries relevant to parental leave taking: Where the boundaries are maintained, where they are challenged, and the role that current technological artifacts play in this. In some settings individuals do not act as tribes; particularly the caseworkers sometimes act as a collective driven by a joint purpose. Nevertheless, the concept of tribes reveals interesting CSCW mechanisms in our particular case. We situate the idea of a web-based, shareable timeline artifact in the analytic findings, and use this to expose the challenges and possibilities of such an idea (Figure 2).



Figure 2. Workshops

Parental leave takers: Continuously changing constellations of peers

Constituent of all parental leave is obviously a child being born. As such, a parental leave case is initiated by the birth of a child. Not the child in itself, but a birth. Once the child, as all Danish citizens, is granted a CPR-number, this number now becomes the object against which information is cross-searched within the municipality. E.g. caseworkers may cross-reference the CPR-number of the father and that of the child to verify that he is in fact the biological father. In contrast to this view of what is constituent to parental leave, the mothers involved in our study gave the clear impression that while the child is of most importance, their maternity leave was a product of negotiation between themselves, their partner and prospective employers. The employers are affected by the leave both by the absence of the parent, and economically by the prospective refund they should receive according to the existing agreement. The time plan can in itself be a complicated matter, as the legislation allows part-time leaves and the scheduling of vacation time intermixed with the leave. Because of the timeframe, parental leave planning sometimes involves the leaves pertaining to several children, and consequently several partners or parents may be involved. It goes without saying that while most couples may have healthy relationships, the sharing of children between broken-up couples can entail challenges in coordination. Even though the configuration of the actors involved stays the same over the potential nine-year duration of the leave process, the boundaries of the family and hence the actors contributing to the planning process may change.

The above-mentioned planning process can be very complex in its constellation of involved actors and is furthermore driven by a difficult balancing of time (spent with the child) and money. Both parameters are central to the

parents' understanding of a sound base for their child. In seeking a solution to this equation, parents in many cases rely on information from a complex network of friends, family, the municipality, their Mothers' group, Internet websites, etc.

When using the services from the municipal office, citizens become members of the tribe of parental leave takers. The tribe that the citizens form with each other, is complex and based on a number of aesthetic and ethical values: The caring of the baby, the distress of being deprived of sleep, etc. Apparently the negative commonness of these tribes is especially forceful in creating a strong feeling of cohesion as numerous quotes in this paper emphasize. This dynamic is related to the way in which the boundaries of a certain community are often consolidated by what *does not* belong; as opposed to which activities or aesthetic value *do* pertain to the community (Bødker, et al., 2003). New parents meet around their shared experience of e.g. sleep deprivation and screaming babies. These are hardships they see themselves enduring in contrast to others, who consequently are not seen as belonging to the tribe. Just as this tribe of parental leave takers overlaps with a number of other values, the citizens are also simultaneously members of many other tribes. In this paper, however, we are concerned with how parents relate to each other in consuming services from the municipality. Quotes 2 and 3 illustrate the emergence of a tribe:

Quote 2: I think that the biggest difficulty was to figure out where I should send my papers, because I do not have a proper employer. I stopped working before my parental leave and went on unemployment subsidiaries. There were so many instances involved, when I was to report all the information. It was very difficult for me to find out how to do that. (Mother, A3)

Quote 3: It can't be right that you have to spend that much time searching for the rules (Mother, A3)

These two quotes exemplify a collective emotion of *difficulty* or hardship within the tribe. This emotion is the most prevalent within this tribe. There is a rational component of experiencing difficulty—searching for the right solution and making errors along the way due to complexity. However, what we wish to depict here lies closer to the emotion of almost biting the dust or in being the loser in a game you do not really understand. This is an experience not uncommon to many citizens in our case context. These new parents see themselves in contrast to the caseworkers of the municipal office who, in their view, should know the rules and possesses an overview. Despite the fact that the employers of the respective parents and their unions are equal sources of rules (and thus potentially, frustration), they are *not* the targets of the frustration of the parents in the same way as the municipal office. The human face of the municipal office, at the boundary between the municipal office and the parents, seems to have a paradoxical role in the communication and coordination between parents and the municipality. At times it seems to trigger the citizens' frustration; while at the same time many of the new parents find the face-to-face (or telephone) meeting less essential to the communication than do the caseworkers.

The citizens' sense of difficulty is used to legitimize behavior that might be deemed immoral under other circumstances. In other words, the tribe develops a set of ethics.

Quote 4: If I was told I had made a mistake, I would probably think, why didn't they make some simpler rules, then? (Mother, A3)

The parental leave takers do not feel that they can actually be held responsible. They have a sense of being a community, which does not have to be fully in control of what is happening. Although it is not a pleasant feeling, it also indicates a basic trust in the system. The parental leave takers have in a sense renounced their being fully competent in handling the process. This sometimes shows in their negotiations with caseworkers (Quote 5):

Quote 5: I had a woman calling; she was quite upset. She couldn't see why she had to go and read law books to get answers to her questions regarding her rights (Caseworker, A6)

Furthermore, although we did not encounter any outright cheating, not everybody felt obliged to tell the caseworkers everything about their situation:

Quote 6: I spent a lot of time trying to understand the rules in the beginning. (...) You probably shouldn't give them too much information. What if they tell you: "You know what, then you are not getting any money?" (Mother, A3)

To keep some information secret is not a code of ethics that comes from some opportunistic general rule of conduct. It is rather a result of being unable to understand the consequences of the legislation when applied to one's own situation, because of the aforementioned *difficulty*.

It is very difficult for parents to disentangle and overview the overlapping sets of rules and regulations governing parental leave, as illustrated by Quote 5. Essentially parents are indifferent as to whether the leave funding originates from their employer or from the state, as long as they get the money they are entitled to, and as long as the plans they make are legitimate.

The citizens engage in tribal communion when meeting with each other and recognizing themselves as parental leave takers. The tribe, however, does not visit the municipality office together, and in this sense the tribe does not do much to expose its customs to neither itself nor the outside world. In this sense, parental leave takers are not a strong community with a clearly defined set of values, but loosely structured heterogeneous tribes of peers continuously proliferating, changing and dissolving as the involved actors' paths cross.

Caseworkers and the municipal office

Much like the parents, the caseworkers of the municipal offices share a common ethical and aesthetical understanding of their purpose and demarcation in relation to the parents to whom they provide service. The caseworkers and their middle managers quite clearly expressed tribal values through various interviews and workshops. They balanced their everyday activities among themselves and in their encounters with citizens between service on the one hand, and control on the

other. It became clear that the caseworkers often operated with the assumption that the citizens at some point needed to meet a human face in their contact with the municipality:

Quote 7: There is a lot of psychology regarding money. (...) People like to have confirmed by a person that they are doing the right thing (Caseworker, A7)

Interviews with many caseworkers illustrated how, during phone calls, the caseworkers tried to personally offer as much information as possible. They do not have e.g. a webpage or a FAQ to point to, nor do they guide citizens on to websites where parents discuss parental leave among themselves (e.g. www.navlestrengen.dk). In one of our prototyping workshops, the researcher presenting the prototype pointed to the fact that the caseworkers could tell parents to call back when they had experimented with the timeline on their own. Still, however, the caseworkers took control of the timeline and started adjusting the parental leave proposal. The caseworkers seemed to find direct interaction with citizens' imperative. Consequently, they focused on personal guidance and advice as central elements of good service, as opposed to e.g. advising the citizens on how to explore possibilities on their own. As the above quote indicates, many of the caseworkers believed that this personal dialogue was pivotal in providing good service and creating trust between the two parties. There were, however, still clear limits to the service provided by the caseworkers:

Quote 8: If she is unemployed she has to go to the people who have the information—her employer or her union. If she has no union she is in trouble. (...) If she has no union and no employer she cannot proclaim that she has a good paternity leave agreement! (Caseworker, about a scenario, A6)

The municipal caseworkers did not see it as their role to advise parents regarding the collective agreements; neither did they have any means of actually keeping an overview of these many agreements themselves. The overall national legislation on parental leave is in itself complex, and some caseworkers even kept from counseling on this legislation as quote 9 indicates. This was indeed an area where the limits of both the tribal responsibilities and boundaries of good service were continuously articulated. In current praxis, they are also clear: the labor unions and many employers have informally accepted the task to explain the collective agreements when a related citizen seeks their advice. Concurrent with this demarcation of responsibility the caseworkers and the middle management often pointed to the office's fundamental responsibility of controlling that citizens complied with the law.

Quote 9: If she (a persona) is unemployed, she has to see a lawyer about her rights (Middle manager, municipal office in workshop, A6)

It was not only in regards to these aesthetic considerations that the caseworkers addressed the relationship of their tribe with that of the citizens'. In regard to ethics, they were quite openly concerned with reducing their own workload. At times this implied avoiding acts, which would entail more questions or demand for better service from the citizens. This was a result of a complex consideration

of a number of concerns. These included a homogeneous service level for all citizens, and balancing the overall work and cash flow both within the organization and from the state to the citizens.

Having outlined the form and context of citizen and caseworker tribes we turn to actual and exploratory boundary negotiating artifacts, with a view to the future.

The back-office system

There are many ways in which the current technology enforces the boundaries of the office in relation to the citizens. The main system is essentially not made to capture the flexibility of the current legislation. Consequently, when citizens wished to make use of the flexible legislation this often induced a number of case processing tasks within the municipal office. If a citizen coordinated an unevenly distribution of leave days with her employer (e.g. every 5th Monday) a caseworker continuously had to manually start and stop the leave case within the system. The solution was often that the citizen sent emails before every leave day, and/or the caseworker would set reminders within the system to start or stop the case again. The caseworkers accordingly called these situations *start-stop-cases*.

Obviously, a very *inflexible* system combined with a *very flexible* legislation is a problem in itself. The caseworkers worked around this lack of flexibility, based on email messages/calls from the citizens and reminders in the system. This somewhat ad hoc solution had become the main artifact on the boundary between parents in such situations, and the back-office system, which ensures that citizens get their money. At the same time, the municipal organization holds on to this back-office system, because it quite efficiently helps them control and execute the cash flow between the many parties involved: Parents, employers, etc.

The example of the start-stop-cases illustrates how the caseworkers spend large amounts of energy maintaining the boundaries. They do this mainly to support the parents in being granted their legal rights, thus shielding them from the evident problems of their back-office systems. Caseworkers draw these boundaries, fragmenting the exchange in order to reduce complexity. This is an important reason why future technologies do not simply need to grant citizens access to improved versions of the back-office systems. However dysfunctional these back-office system may seem, caseworkers need to handle many complex behind-the-scene matters, which inherently call for more complex maneuvers than performed by citizens. This dilemma challenges a dominating concept within the current e-governance and citizen service discourses: That caseworkers and citizens should use the same web-based solutions (KL, 2006, p. 8).

With this perspective, the timeline idea becomes a thinking tool, not just for us. It challenges social and collaborative web technologies: Our timeline design idea is not a shared web system that replaces existing back-office systems and it is not obvious that they could or should be. Neither is it a matter of simply sharing

information or FAQs. It is inherently collaborative and shared between several users; a simple shared object, which can be seen and handled by both citizens and caseworkers, together or apart.

Information or control

Quote 10: It is fine with good citizen service, but our hidden agenda is to make it easy for us (...) and to push some further service or information to the citizens will mean a clash of these agendas. (Middle manager A6).

The tension between better service and the workload was mentioned at several occasions by the caseworkers. Moreover the caseworkers pointed to the inherent paradox of being both a control agency and a supplier of knowledge and service. Accordingly caseworkers as well as their managers voiced the concern that if they pushed more information, e.g. referred to a FAQ, it might lead to more questions.

In other words, they worried that such information might open their boundary for what could be seen as further penetration. The above analysis provides a different perspective: The expecting parents seemed interested in clear up-front information, which could qualify and limit their questions to the municipality. The caseworkers were, however, apparently more focused on balancing their workload and as such maintaining their boundary to the citizens, as opposed to providing quality service at possible increased costs. However, one caseworker voiced a different reason for doubting the result of citizen-to-citizen services:

Quote 11: If I were a citizen and I went to a site where the municipal office supported citizens in sharing case stories and timelines, perhaps I would think—what is this really? Is it just a way for the municipal office to shortcut questions and make citizens do parts of their work? So the office makes it easy on itself, letting me chat with my friends instead of getting proper advice. (Caseworker, A6)

The caseworkers were generally interested in exploring ways of providing information to citizens, while simultaneously reducing their own workload. In a workshop, caseworkers vividly discussed whether the doctors' or midwives' offices might be better locations for information (pamphlets, on-site computers, etc.) about parental leave. Moving the challenging information obligation away from the office would be a way of clearly drawing the boundaries of the municipal office, allowing for a more direct emphasis on the office's control function. Such a restructuring of boundaries does, however, beg the question "who's responsibility is it then to provide such information?"

Our explorative timeline design challenges current boundaries. It allows for new parties to contribute information and solutions. It is dependent on the sharing of plans, and on various actors being able to validate and provide parts and packages of rules and information. This will inevitably call for a renegotiation of current boundaries for information and control.

Meeting places for the tribe of new parents

New parents typically communicate with the municipal office individually or with their spouse. In this manner, the tribal aspects of parental leave are quite absent from the interaction between the municipal office and the citizen. As such there is no strong sense of community among parents in their present interaction with caseworkers. The tribal aspects are, however, widely voiced where parents meet. The Mothers' groups and networks of friends are particularly often sources of the sharing of experiences with parental leave planning. Here parents recapitulate the hardship and confusion, which the complex constraints often entail. This leaves many of the discussions and recommendations open and ambiguous. As a consequence, advice on how to softly stretch the legislation is quite normal, leaving solutions unsuitable for direct exchange with the municipal office (quote 8).

Parents, in addition, meet through various websites, e.g. Navlestrengen.dk, but these meeting places are in no way connected to the municipality. Although the municipalities as a whole have the richest know-how concerning parental leave planning, there is currently room for actors outside the municipality to create web solutions, which provide citizens with the possibility of feeling they belong e.g. to a parental leave tribe. Such feelings of belonging are neither inherently good, nor bad to the municipal office and the parents. The caseworkers in general know little about such sources of information, and they hesitate to refer to them in their interaction with parents. The caseworkers' main concern is that incorrect information would come from the municipal office. Since the mentioned websites are outside the boundaries of their control, referring to them may jeopardize trust in the information originating from the office. Consequently, caseworkers prefer to draw an explicit boundary between such websites and their own information.

Based on the contrasting perspectives of the parents and the municipal office, a shared social network-meeting place such as a Facebook group does not seem appropriate. It is more likely that parents wish to separate their communication with each other from that of their interaction with the municipal office. Likewise, the municipal office has good reasons for not interfering with parental leave tribes, while targeting other ways of improving the information regarding parental leave.

The timeline design interestingly exposed how issues of sharing may also involve sharing without assessing the quality of the content. Both the paradox of the municipal office's obligation to simultaneously control and provide service to citizens, as well as their fear of taking responsibility for incorrect information originating as a consequence of participatory citizenship, potentially call for a reevaluation of the current legislation.

Sandbox mechanisms

When presented to the timeline in workshops, caseworkers saw this new kind of artifact as a way of obtaining a better overview. They were concerned with how

the rules would be enforced in the timeline, mainly to make sure that they would not have to do such checking while communicating with parents. The parents in a similar workshop focused less on the artifact as a facilitator of the communication with the municipal office, and more on the artifacts potential as a tool for experimentation in the form of a sandbox. This function allows for personal exploration of the rules and agreements, and generation of the needed documents for their various employers, unions and the municipal office.

This exploration of rules and finances can take place before the citizen turns to the municipality or their union. Hence the solution can actually reduce the amount of calls to the municipal office. The idea is that the timeline shows the documents needed for communication between the parents and the municipality or other agencies. Simultaneously it exposes and prompts for the time constraints and deadlines related to the documents. The information in these documents may be generated semi-automatically based on the timeline, but they will only be sent to the designated receiver when explicitly confirmed by the citizen. The timeline may only be shared under similar circumstances.

The timeline accordingly is a new artifact on the boundaries between the municipal office and the citizens. It allows citizens to maintain and control the boundary, while exploring and experimenting in private—alone or with their spouse (Grudin, 2002). In this sense citizens can control and inspect the totality of the information sent to the municipal office.

In summary, the sandbox mechanisms underpin one type of user's exploration of a set of rules (outlined by another type of users). This exploration is left open and uncommitted, until parents explicitly decide to share the information with caseworkers crossing the boundaries. Current web designs to a large extent leave such boundaries unclear. Information entered into web forms can only seldom be saved for personal use or later revision. Our aim is different.

New participants

A timeline-planning tool adds value to both citizens and caseworkers. Even more so if e.g. labor unions or major employers become active participants, providing their regulations and agreements as forms or plug-ins to the timeline. This could potentially fill the gap between the lacking knowledge of both citizens and caseworkers concerning collective agreements. The municipality ought not be responsible for such information. This would move current boundaries making employers and labor unions responsible for providing this information. At the same time it would, however, make the total set of rules and regulations more visible to citizens, and relieve the municipal office from their commitment to informing each and every citizen about the rules that apply to them. It would also address the individual parents' need to overview their salary according to collective agree-

ments and subsidiaries paid by the municipality. This was a need that was strongly indicated in workshop A3:

Quote 12: There should be a genius somewhere, who gathered all the rules at one website – and that page would be all that was needed (Mother, A3).

A shared timeline with plug-ins would provide the foundation for the work of such a “genius” who might proliferate from within the tribe of parental leave takers. However, it is a very open question still how the municipal office as well as other agencies would take responsibility and guarantee the quality and validity of the timeline elements. It seems that clearly defined plug-in components draw boundaries that render visible the connection between information, constraints and particular communities. Who should be responsible for the maintenance of such components is a question that still needs answering.

Moreover a web-based timeline makes it possible for friends to share parental leave plans. It is even possible that a powerful agent will gather, “all there is to know about parental leave” on the Internet, for citizens to use—possibly at a cost. A very relevant question would then be, if this will and can be done with the cooperation of the municipal office?

What is needed clearly differs from open communities sharing simple information, person to person. Boundaries may be altered, but the main goal should not be to remove these borders. Rather, new forms of boundary control between tribes and communities should be put in place and these new mechanisms of control should be visible to all users.

Participatory citizenship beyond Web 2.0?

Introducing a collaborative technology allowing contributions from both citizens, municipal caseworkers and e.g. labor unions severely challenges the caseworkers’ understanding of their own role towards the citizens. It raises both legal and ethical issues in relation to the municipality’s responsibility of validating such information. Furthermore, it opens up for new constellations of collaboration and communication allowing citizens to help each other, and e.g. labor unions to provide easily accessible additional information.

Both of these possibilities are, however, dependent on a somewhat philanthropic *participatory citizenship*, where citizens spend time helping fellow leave-takers, and labor unions provide e.g. plug-ins to the web service. Following this line of thought, citizens will accordingly have to change their perception of boundaries and the responsibility within the division of labor between themselves and the municipal office. To a larger extent some citizens have to master the complex legislation if they are to be able to provide consistent answers to their kinsmen. This calls for personal experimentation in exploring the possibilities and limitations of the legislation, and for a move away from what we identified as a common experience of difficulty and frustration.

There are many reasons for maintaining the boundaries between citizens on the one hand, and the municipal office on the other. Several back-office systems have functionality for securing safe case-flows, cross-referencing that only entitled citizens get financial assistance, etc. These functions are a prerequisite for well-functioning e-governance but they are complex. Furthermore, judicial functions cannot easily be transferred to those getting the financial assistance. Caseworkers have to use these systems whereas citizens do not. This reinforces the boundaries between the two tribes, and makes fragmented exchange across a clear boundary important. Nonetheless, new technology at the boundaries may change the aesthetics of both tribes, and e.g. the human face at the boundaries may lose its appeal (to the caseworkers). Such changes might in fact be required if this new technology is to function well. Alternatively, the municipality will have to be kept entirely out of the loop in the development of technologies that allow for participatory citizenship.

Participation may emerge from ethically driven solidarity with other tribe members, and other forms of community mechanisms. Boundaries need to be changed, and in some instances enforced to make this happen. New artifacts on these boundaries may facilitate such changes. Here well-known issues such as trust and openness (e.g. Bannon & Bødker, 1997) call for attention. Trust is at present solely an issue within the direct interaction between the citizen and the municipalities. Other actors may, however, take on new responsibility, e.g. by providing information about the legislation, and both municipality and citizens need to reconsider divisions of responsibility and trust accordingly. Citizens' private experimentation and controlled sharing behind certain boundaries are similar challenges. In this fragmented exchange, caseworkers are spared from issues concerning beliefs in upbringing, family values and other concerns beyond their professional competencies. In their present form most municipal documents and web forms are only used to facilitate a direct information flow between citizens and the municipality. They are neither instruments for the citizens to explore the rules, nor for them to e.g. save information without sending this to the municipality. With a more multifaceted use across tribal boundaries, it seems pivotal that all actors should be able to see and control the boundaries of openness.

Accordingly, a Web 2.0 solution is not just an opening up of communities. Some boundaries need to be maintained and changed as well. A sandbox function enables experimentation in private, while the possibility of sharing opens for collaboration with caseworkers, as well as facilitates a participatory information flow between citizens and within the smaller tribes of Mothers' groups. Altering of boundaries is a complex task involving several stakeholders such as legislators, caseworkers, municipal managers, citizens, and developers.

In order to address the challenges and possibilities we see for social and collaborative technologies at large, we take the concept of Web 2.0 as a starting point. Web 2.0 is based on governing principles such as network effects, where

the quality of the technology increases with more users and continuous development with/by the users, adapted to their behavior (O'Reilly, 2005). The current state of Web 2.0 is that it takes many interesting forms in voluntary communities, but remains to be explored in settings where purpose and efficiency play important roles. The presented case, and the eGov+ project attempt to do exactly this. The case of parental leave, however, also demonstrates some of the major challenges for the Web 2.0 paradigms: Openness is not a one on one exchange of data, and service between individuals does not necessarily imply one big, happy community. We propose that a next important step in moving beyond Web 2.0 is to consider more than individual actors in one voluntary community. Tribes and boundaries need to be recognized and explicitly addressed in design and use.

Tribes as a way of providing a less purpose-driven focus on communities and boundaries has been useful in discussing the community of expecting parents, whereas the caseworkers comply better with more purpose-oriented definitions. Tribes helped analyze the heterogeneous groups of citizens, although we did not identify all groups in play within parental leave. The tribal perspective revealed the naivety of constructing Web 2.0 solutions where both caseworkers and citizens have, or need to have, a strong, shared sense of community. Thus, the concept of tribes proved its merits in addressing boundary negotiation.

We agree with Lee (2005) in seeing boundary negotiating artifacts as a strong concept, analytically as well as design-wise. A strict focus on issues within a boundary, or status quo of such boundaries is likely to result in overlooking important design potentials. In line with Bødker et al. (2003), we propose the turn towards dynamics of boundaries, boundary maintenance and boundary change instead. In doing this we have identified interesting alternatives to current practice.

Acknowledgments

We wish to thank the eGov+ team and The Danish Strategic Research Council as well as the involved caseworkers and managers and Dorthe Haagen Nielsen.

References

- Bannon, L. & Bødker, S. (1997). 'Constructing Common Information Spaces'. In Hughes, J., Prinz, W., Rodden, T. & Schmidt, K. (Eds.). *Proceedings of ECSCW97*, Dordrecht: Kluwer, pp. 81-96.
- Barth, F. (1969). *Ethnic Groups and Boundaries*. Boston: Little, Brown.
- Bias, R.G. (1994). 'Pluralistic usability walkthrough: coordinated empathies'. In Nielsen, J. & Mack, R.L. (Eds.). *Usability Inspection Methods*. New York, NY: Wiley, pp. 63-76.

- Bødker, S. (2004). Status, tax authorities and citizen services in Aarhus. Status report, 2004.
- Bødker, S., Kristensen, J. F., Nielsen, C. & Sperschneider, W. (2003). 'Technology for Boundaries'. In *Proceedings of GROUP'03*, New York, NY: ACM Press, pp. 311–320.
- Bødker, S.(1999). *Computer applications as mediators of design and use—a developmental perspective*. Doctoral dissertation, Department of Computer Science, University of Aarhus, DAIMI PB-542.
- Bohøj, M. & Bouvin, N.O. (2009). 'Interacting with casework documents using time lines', in press for *Hypertext 2009*.
- Clement, A. & Wagner, I. (1995). 'Fragmented Exchange: Disarticulation and the need for regionalized communication spaces'. In Marmolin, H., Sundblad, Y. & Schmidt, K. *Proceedings of ECSCW'95*. pp.33-49. Dordrecht: Kluwer.
- Gaver, W. W., Boucher, A., Pennington, S., & Walker, B. (2004). 'Cultural probes and the value of uncertainty'. *interactions* 11, 5, 53-56.
- Greenbaum, J. & Kyng, M. (Eds.) (1991). *Design at Work: Cooperative Design of Computer Systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Grudin, J. (2002). 'Group Dynamics and Ubiquitous Computing'. *CACM* 45(12), 74-78.
- KL (2006). *Kommunernes digitale strategi*, <http://www.kl.dk> (Downloaded 03032009)
- Lee, C. (2005). 'Between Chaos and routine: Boundary negotiating artifacts in collaboration'. In Gellersen, H., Schmidt, K., Beaudouin-Lafon, M. & Mackay, W. (Eds.), *Proceedings of ECSCW 2005*, pp. 387-406.
- Luitjens, S. (2008). *Learning by comparison. E-Government in the Netherlands*, oral presentation at Forvaltning & Digitalisering, (<http://cok.dk/default.asp?id=76458621810573504012008>).
- Maffesoli, Michel (1996). *The Time of the Tribes: The Decline of Individualism in Mass Society*. London: Sage.
- March, J. & Simon, H (1958). *Organisations*. New York, NY: John Wiley.
- O'Reilly, T. (2005). 'What Is Web 2.0'. O'Reilly Network.
<http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>.
- Plaisant, C., Milash, B., Rose, A., Widoff, S., & Schneiderman, B. (1996). 'Lifelines: Visualizing personal histories'. In *Proceedings of CHI'96*, New York, NY: ACM Press, pp. 221-227.
- Scaife, M. & Rogers, Y. (1998) 'Kids as informants: telling us what we didn't know or confirming what we knew already?'. In A. Druin (ed.) *The Design of Children's Technology: How We Design, What We Design, and Why*. San Francisco, CA: Morgan Kaufman, chapter 2.
- Star, S.L. (1989). 'The structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem solving'. *Distributed Artificial Intelligence*, volume II, chapter 3, San Francisco, CA: Morgan Kaufmann, pp. 37–54.
- Star, S.L. & Griesemer, J.R. (1989). 'Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's museum of vertebrate zoology'. *Social Studies of Science*, 19, 387–420.
- Wimmer, M.A. (2002). 'Integrated service modeling for one-stop government.' *Electronic Markets*, 12 (3), <http://www.electronicmarkets.org/issues/volume-12/volume-12-issue-3/>.

Research Project as Boundary Object: negotiating the conceptual design of a tool for International Development

Ann Light¹ and Theresa Dirndorfer Anderson²

¹Sheffield Hallam University, ²University of Technology, Sydney

¹*a.light@shu.ac.uk*, ²*theresa.anderson@uts.edu.au*

Abstract. This paper reflects on the relationship between *who* one designs for and *what* one designs in the unstructured space of designing for political change; in particular, for supporting “International Development” with ICT. We look at an interdisciplinary research project with goals and funding, but no clearly defined beneficiary group at start, and how amorphousness contributed to impact. The reported project researched a bridging tool to connect producers with consumers across global contexts and show players in the supply chain and their circumstances. We explore how both the nature of the research and the tool’s function became contested as work progressed. To tell this tale, we invoke the idea of boundary objects and the value of tacking back and forth between elastic meanings of the project’s artefacts and processes. We examine the project’s role in India, Chile and other arenas to draw out ways that it functioned as a catalyst and how absence of committed design choices acted as an unexpected strength in reaching its goals.

Introduction

The paper introduces Fair Tracing, a UK-led interdisciplinary project to research a bridging tool connecting producers with consumers across global contexts. We offer this work in conjunction with the conception of ‘boundary objects’, introduced by Star (eg: Bowker & Star 1999; Star & Griesemer 1989) as a way of acknowledging and exploring trans-disciplinary and trans-cultural interactions. In this way, we intend to show how the openness of a research agenda in combination with the power of some central ideas came to help embed

propositions from the project in contexts of use. In the process, we demonstrate how the tool's function and even the meaning of the design process became contested as work progressed. And we document the many perspectives that arose as the initial team widened to take in local research and business partners in other countries and as different interests offered their positions with regard to developing and using the projected software. In this way, we explore why, instead of narrowing, as most design projects do when research and prototyping start to reveal suitable constraints, the strength of the core idea turned it into an ever-widening site for multiple (often incompatible) versions of a Fair Tracing system.

This paper has two objectives: first, it seeks to contribute to our understanding of working with ICT across cultures and in a "Development" context. Second, it aims to extract more general value for CSCW from looking at an interdisciplinary research project with money and ambitions, but no clearly defined beneficiary group at start, and how its amorphousness contributed to its impact. These objectives recognise that there are many challenges to meet, including that:

- ICT for/and international development (ICTD) projects have a high failure rate in terms of uptake, even when a functional application is developed;
- Many ICTD projects are initiated without the intended beneficiaries; but are conceived by exogenous parties to improve others' wellbeing;
- Societal contexts, unlike workplace productivity contexts, do not constitute themselves into clear user groups for research teams to collaborate with in defining issues, setting boundaries and doing design;
- In Europe, as elsewhere, funded research is being increasingly required to show impact as well as the potential for it.

Before commencing, one ambiguity must be dealt with. The authors have accepted the convention of talking about 'Development' and 'International Development' to refer to promoting socio-political change in relations between global citizens and engaging in knowledge exchange across cultures. Leaving aside the legitimate discussion of whether such 'Development' is possible, desirable or can be supported with ICT, it is worth drawing attention to the convention used here that 'Development' refers to this domain, while, with lower case letters, it is the standard technological use of making products and services.

Project Background

In this section, we present an overview of the history and goals of the Fair Tracing project, the institutions and researchers and how they came to work together.

The Fair Tracing idea

The idea behind the Fair Tracing system is simple and powerful. It is a publicly-available tool that makes the provenance of any goods more visible by illustrating

the supply chain from producer to consumer and, in so doing, pushes practices towards ethical production and purchase. In particular, it is intended to give visibility to small producers in emerging economies operating in global markets. Indeed, the name of the tool acknowledges its relation to the Fair Trade movement which supports producers in developing countries committed to social investment in their communities. The tool is to improve life for both ethically minded traders and consumers who wish to understand and discuss the origin of their purchases.

The idea caught in the Fair Tracing project is immediately comprehensible and has been popular with everyone from designers to producers to consumers. The name ‘fair’ wins instant approval. However, it is obvious there are complexities below the surface, such as, what constitutes ‘fair’: the tool opens up a discussion of design for social justice (Light & Luckin 2008). Practical implementation is non-trivial because of many factors, not least that it incorporates two key human activities (production and consumption), doing so specifically in a contextualised way to allow sameness and difference across the world to be determined. It deals with values and their presentation across cultures; brings issues of information retrieval and architecture; has to accommodate differences in producers’ readiness to contribute information and meta-data about it (which can be automated in IT-mature contexts) so as not to affect the overheads of use beyond tolerance, etc¹.

The project was funded to research the building of such a bridging tool and contribute understanding of its potential for implementation and use *in context*, beyond individual technological components², over a three year period till 2009.

Genesis of the Idea

The Fair Tracing idea was conceived in a workshop on socio-technical ways of “Bridging the Global Digital Divide” (BGDD), run by the UK’s Engineering and Physical Sciences Research Council. Because of this birthright, the thrust of the project was using information technology to put small-scale producers in emerging economies in touch with their markets overseas. Four researchers grouped round this idea in the workshop, later to evolve into the project’s management team. The four comprised a political scientist with links to India (M1); an economic geographer with Chilean experience of participatory practices (M2); a social informatics/HCI specialist concerned with the politics of design (M3) and a computer scientist into security and ethics (M4).

¹ There are other complicating factors, such as the way that products change nature en route, e.g. grapes from multiple vineyards becoming wine, or become other products, like cotton and buttons becoming shirts. There is the sheer number of elements to be recorded and manipulated if every instance of a product is to get a moment of analysis through all stages of production.

² So, for example, it was clearly possible to organise a value chain tracking system using the assignment of numbers and chronology to each producer and production event, but no one knew how far this process would be of use *in situ* along the value chain and for interested third parties.

Driven by the research interests of the team, the project took shape: to provide user-generated economic, environmental and social information about each leg of the supply chain (Porter 1985), showing production practices and actors at every stage of the process as far as purchase and with opportunity for purchasers to add experiences too. Its context was defined as the global network, where the identity management of goods using technology such as RFID could give large companies advantage over others. Thus the tool being researched was conceived to address a growing divide between networked players and those for whom business would become relatively more difficult, not because their production conditions had materially worsened, but because a new competitive practice has been introduced.

The technological focus of the project was determined by the funding. The funders had implicitly decided that a tool would ensue and that, in exploring existing socio-technical systems, workshop participants would be assessing how to make interventions. The multidisciplinary team forming held a range of views on this: some more interested in research overview, some keen to show a gadget and initiate change; some functioning at an abstract level and some keener to understand a workable assembly of interactions. After debate, it was agreed that the tool would be conceived as generic, “Open Source” and supported by an infrastructure that would make it robust and cheap enough for everyone to use. It would offer low bandwidth multimedia for sending stories and facts attached to individual instances of items. It was recognised that it would not be viable to build and test such a tool as part of the project: not only was the project short on development time, but as research it would be unethical to put an end-to-end prototype into small producers’ infrastructure. Instead the team decided to build relations with producers as partners: bringing them in as informants to a theoretical investigation to consult on ideas, prototypes and uses.

Key Players

When we refer to ‘participants’ in a research project, we often mean people who are brought into the process for the purpose of collaborating upon or evaluating an artefact’s design. Because we are looking at the research project as an artefact in this paper, we view all players (the researchers, funders, intended beneficiaries) as participants creating the research process and determining the structure.

In addition to the management team of the four researchers (M1-4) mentioned above, the project committed to working with two case study partners (to include small-scale producers and their contacts along the chain), and, further, to fund:

- Two PhD students with CS and IT/HCI backgrounds (PhD1 and 2);
- Other students building interfaces based on the research data for projects;
- A local social researcher at each site of the case studies (LSR1 and 2);
- A research assistant, initially drawn from the social sciences (RA1) and subsequently (covering maternity) bringing interaction design skills (RA2).

Subgroups met and collaborated as follows: management overview: (M1, M2); research overview: (M2, M3, RA1/2, PhD1); technical overview: (M4, M3); infrastructure: (M4, PhD2); design: (RA2, M3, [M2]); Chile: (M2, M3, LSR1); India: (PhD1, M3, [LSR2]); consumer study: (M2, M3, RA1/2). Responsibility was devolved within these subgroups and it is indicative that the team rarely spoke through one mouthpiece when together but kept the distributed feel of the project by presenting the research serially to the BGDD network and others.

Extending the team, the management subgroup chose its case study partners pragmatically, reusing existing cultural knowledge and relationships. This resulted in pursuit of partners in Chile in the Fair Trade wine industry and in India to work with the shade-grown coffee of Karnataka. However, there was a resulting lack of symmetry in partners and dynamics with them. The value chain of Chilean Fair Trade wine is straightforward to follow and has few steps³. Not so that of Indian coffee: only some coffee estates in India are marketed discretely and these are not the small ones; none have Fair Trade certification; the beans of the smaller producers lose their identity in an anonymous pile at the curing works where they go to be graded; quality crops and those with an ethical story are buried along with inferior beans; multiple traders get involved. The team could choose between following a traceable coffee line, or staying true to working with small producers - and untraceable goods. The team chose the latter, partly to increase the design space investigated. That said, the political nature of the Indian coffee sector intruded so there was no clear boundary about who was involved: arguably the whole Karnataka coffee industry.

Decisions such as these were negotiated at face-to-face team meetings, which routinely excluded, because of distance, some members of the wider team who needed information. Augmenting this stuttering flow of information, a broader email list received a weekly update of activities through the central point of the project manager (M2). Supplementing the email stream was a shared project blog.

Scope and Method

Fair Tracing's studies lasted three years and were distributed across a wide set of possible enquiries. Specifically, the team set objectives to gather understandings of the value chain actors about the chain, to learn about actors' working practices and use of ICT, to explore the needs and desires of the different chain actors and what they thought a Fair Tracing type tool would be and do. The team also sought to establish how much work actors would put into to developing stories for it, inputting data into it and preparing their material for the different cultural contexts it would be shown in – since it was conceived to be a Web2.0 style tool.

³ Although the politics is never simple: when first approached the subsequent partners referred the FT team up the chain to their importers for approval before agreeing to collaborate, who in turn required permission from their distributors, a major supermarket chain. Once the supermarket had agreed, each previous stage accepted the collaboration and a collaboration agreement was drawn up.

Overall, they hoped to gain an overview of what the two case study chains wanted, as far as the consumers and to do so in a collaborative way, with long-term partners so that relationships could be established. A flavour of the research is provided here. The detail is outside the focus of the paper (but see, for instance, Light et al 2009).

Research activity: what the subgroups did

At the producer end, the relevant subgroups built a relationship with the partners in Chile and India and visited them to conduct together:

- scoping of ICT use, views and knowledge of value chain, collaborative identification of key elements of production processes to communicate, willingness of local people to collect and share stories, and design workshops to explore turning knowledge into representational material,
- reporting on the project, testing ideas from consumers, evaluating changes in interpretation/desire for a tool, trying out contrasting prototype interfaces as an elicitation technique to learn more about information to share, how and why.

The research subgroup then mapped the use of relevant tools and implications this had for input of data, while the consumer subgroup concurrently conducted:

- accompanied shopping trips with a diverse range of British consumers,
- a survey of shopping priorities,
- interviews with consumers in front of interfaces, representing spatial, temporal and social metaphors for the presentation of value chain information.

So extensive investigative work was undertaken and, in addition, the wider team:

- explored IT platforms for peer-to-peer robust storage,
- built prototype interfaces for different platforms, such as Web and iPhone,
- analysed the complexity of the value chains and their implications,
- ran a seminar for other organisations interested in tracing technologies,
- wrote the blog and gave a number of talks, largely in the UK and Europe.

Most data were collected through semi structured interviewing and also some limited ethnography with the business partners. Most exchanges were recorded in audio files, though some were videoed and some sensitive speculative meetings were only recorded in written note form by the members of the team present. The research subgroup also made records of planning conversations and meetings.

Now that we have presented the project, we will situate our discussion by looking at the literatures on Development, participatory design and designing for appropriation and use these as a way of teasing out some issues facing the project.

Processes and Principles

There is a long history of participatory design (PD) in developing ICT, and, although many of these projects address tightly defined workplace problems (Muller 2002), there is a political sensibility to involving potential beneficiaries in the design of their tools (Greenbaum & King 1991) in sympathy with the

inspiration for collaboration at the heart of the Fair Tracing research process. In spite of shared history, the nature of beneficiary participation varies and can mean co-designing research structure and defining challenges or can mean helping with predefined goals set by others. To complicate this, participatory practices in ICT are not well established for contexts without organisation and structure (Muller 2002, Irani et al 2009). The PD movement comes out of trade unionism and workplace automation, from socially and technologically coherent environments. PD projects have tended to operate in contexts where there is collective representation for local experts and where defined professional activities exist to address (though see Dearden & Rizvi 2008 on PD in a Development context).

“Before designers can solve a problem, they first must define what it is. How do designers of new technologies begin when they are unsure of what they are making, what it should do, or who will use it?” asks Erickson (1995). This uncertainty must underpin any new research project to initiate a design, but is particularly apparent with projects that seek to serve the wellbeing of others. The next two sections consider the particular challenges Development projects bring, and how these relate to design practices and beneficiary engagement, so that we can draw out how these issues impacted in the Fair Tracing project.

“Development” and designing

ICTD (or ICT for/and [international] Development) projects bring all the challenges of implementing any new system but add a few of their own: they tend to involve working across cultures, often with limited access between designers and intended beneficiaries, and use specialists from more industrialised countries to do design on behalf of communities whose understanding of technology may be very different from theirs (Irani et al, 2009). As Donaldson observes: “Remote design (design from afar) and parachute design (design from afar with visits) do not lend well to capacity building, let alone product sustainability.” (2008:37). Just as the Fair Tracing project was conceived in the Global North, with exogenous research team and funding, so many Development projects come from outside the region to be “developed”. Worse, the systems to be designed often live only in the heads of funder and funded party and have no mapping to existing needs, processes and organisations. Nonetheless, there is the imposition of specifics from funding regimes that need assurances before they support work.

At best, both funding and design team are flexible and will find emergent solutions to ‘problems’ or ‘opportunities’ with locals. However, without a clear idea of either intended beneficiaries or the function of the tool being funded, there is potential to loop. How you cast your net for participants will determine what you build, but shifting your gaze slightly might introduce quite another set of beneficiaries with a different problem to solve or opportunity to exploit. The ‘who’ and the ‘what’ become shifting sands in the project, with no means to prioritise, since local knowledge of conditions is greater than that of the incoming designers. One way of dealing with this exogenously is to move slowly and engage extensively, but funding does not often support this approach.

The Fair Tracing proposition is unusual in Development contexts: it was not predicated on the idea of taking technology to aid a community, but implicates the many ends of an international 'bridge' in choosing to learn new things from each other. It is the primacy of this bridging role for the tool that first calls to mind the notion of the boundary object, a device that (conceptually at least) is rigid enough to be identifiable across contexts, and yet plastic enough to take on a local role and context-specific meaning. We will return to this point later to discuss the bridging function of some of the other boundary objects identified in the project.

A big challenge for Fair Tracing, with the world of producers and consumers as its users, was identifying the stakeholders that own the problem being defined. The partners' view of project activities would be highly determining – in that their thoughts would form the basis of everyone's understanding of the issues – but were also less informed by research into the behaviour of the British consumer (the project's defined end point), than the project team's. This gave the team responsibility as *project* owners to share insights from both ends of the bridge and not give too much authority away to any single group.

The team's authority as informants emerged in another Development context. One justification for introducing specific functions from outside is to assert that there are types of tool available in highly industrialised countries which might be of use in less developed regions where they are not available in suitable form. This position of helping others with their "deficiencies" can be justified by what Sterling (2005) calls the 'Line of No Return', past which a technoculture cannot voluntarily return to an earlier technocultural condition, and the 'Line of Empire', where 'the imperial technoculture can spew its objects and processes abroad, more or less at will' and 'those who lack the productive capacity are forced into colonial or defensive positions' (2005:10). Taking this perspective, it is an act of assertion for colonised regions to develop their own forms of artefacts and processes to level the playing field. If enough of the conception and embedding of these artefacts and processes can be handled by the intended beneficiaries, in theory at least, it is possible they can also avoid being colonised by the *values* of dominating technologies – although this is a particularly fraught argument in the context of commerce and international capitalism. Fair Tracing was predicated on this type of position, in that the identity management of products is joining a worldwide network of information and control strategies, and the system would be a tool with an identity management function for small traders to adopt, adapt and use, either to compete with or to plug into new systems of accountability.

Underlying other concerns in Development is always the question of values. Designing and installing a new system and convincing people to use it is, in itself, a huge undertaking. Yet the act of enhancing wellbeing cannot be expressed in terms of equipping people with levelling tools. These are merely the means to an end. Wellbeing must be expressed in first order terms (Holden 2004) and who chooses what that definition will be? This raises one last Development issue to discuss. Since the proposed tool is a link between different worlds, it not only has a bridging function, but inevitably also a representational function. Representation across global divides and cultures is profoundly political. And not only are there politics in representing a socio-technical system across players, but

the act of doing so has within it the making of changes to that system (see Light 2008). So, who decides on the representation(s)? How far is the system open enough to allow the matter to be decided by actors themselves in use-time, rather than being a legacy of the designers' vision? And, without a single editorial voice, how do multiple different actors establish what is in their best interests as a representation for marketing themselves as ethical, not least as ethical trends change repeatedly?⁴

Participation, openness and indeterminacy in designing

In the section above, we observe how sensitivity to Development's particular challenges, such as the desire to avoid alienating local people, can result in open, ill-defined starting points for projects and varied and under-defined contexts of use. And we have noted that every design has elements of uncertainty in it, certainly at outset. Dorst suggests design is complicated because one partly creates the landscape one will travel through (2003). Fischer asserts that, among other factors, '[c]omplexity in design arises from the need to synthesize stakeholders' different perspectives of a problem' (2007). Drawing on the work of Rittel & Webber (1984), who point out that one cannot understand a "problem" without having a concept of the solution, Fischer (2007) also observes that emphasizing the integration of problem framing and problem solving casts design as a search for a problem space rather than just within a problem space. Further, this emphasizes the importance of problem owners (for whom an artefact is designed) as stakeholders in the design process because they have the authority and knowledge to reframe the problem as the problem space is understood (2007).

For a system to work across multiple or under-defined contexts, it must be open to adaptation. Designing for end-use appropriation, or beyond that, end-user customisation, requires an underdetermined artefact. Fischer suggests a: 'primary challenge of underdesign is in developing environments and not the solutions, allowing [problem owners] at use time to create solutions themselves. This can be done by providing a context and an interpretive background against which situated cases coming up later can be interpreted' (2003). In this way, he unites a participatory philosophy with a pragmatic response in the spirit of distributed use.

If we look at the Fair Tracing system, it is a tool that potentially embeds completely into the context of production, coping with the very different priorities and practices worldwide so that items can be tracked from source to destination. This kind of tool must get out of the way of its users rather than dictate behaviour (a medium rather than a mechanism, in Bentley & Dourish's terms, 1995). The Fair Tracing tool can be seen as something of a hybrid then: with many highly situated voices engaged in problem definition, design and use, on the one hand, and the need for a wide open system for local appropriation as an outcome, on the other. As noted earlier, no single group had sufficient knowledge to design for engagement with others and the project team took the role of mediator but had no

⁴ FT got caught out in an early prototype with a map interface by the sudden emphasis on food/fuel miles at the expense of other ways of recording environmental impact (see Light 2008).

clear mechanism to prioritise functions or representations. All of which left the research wide open.

Discussion of openness in this literature has been concerned with the resulting technological artefacts and not so much the design approaches for making them. In another set of discourses, anthropologists and information science researchers – to name but two examples - have reflected on the interpretive flexibility of a wider range of design phenomena. In the social sciences, the work of Denzin (1989) and Lincoln & Guba (1989) has been influential with regard to taking a more interpretive, emergent approach to research. Design discussion framed in this social science discourse has been shown to value the open, indeterminate, interpretive approach similar to that presented here. In the information sciences, the work of design anthropologists (e.g. Nardi & O’Day 1999, Suchman 1987) has further influenced researchers to develop interpretive, flexible approaches to reveal a richer sense of the socio-technical context. Within design practice, the closest one comes to the Fair Tracing research style is the use of cultural probes, sent out as a research tool to enquire into local meanings (Gaver et al 2004); Critical Design which, like art, serves to challenge people’s perceptions rather than be used in the form conceived (Dunne & Raby 1999); and designs that trade on their ambiguity as a way of investigating potential use (Gaver et al 2004). These have in common with Fair Tracing their desire to solicit multiple and contrasting understandings rather than pin down a single designable angle. However, they have not been focused on the Development context and all have been conceived as an elicitation stage, informing designers through the involvement of users, rather than as a means of engaging user-designers in building their version of the self-same project.

Boundary Objects at Play

To sum up, when we look at Fair Tracing we are examining the interpretive flexibility of an idea. This is distinct from looking at flexibility in implementation of that idea – i.e. in the tool - or using instances of that idea as a probe for informing on other ideas, as is more common in design. The idea has become the design artefact. We now look at how the idea - and various other aspects of the project, such as the metaphor of the value chain - became situated and crossed boundaries. To do this, we first outline the concept of the boundary object.

Boundary Object: concept defined

Star (Star & Griesemer 1989) introduced the boundary object to explain objects inhabiting multiple contexts at the same time whilst having both local and shared meaning. The notion is further explored in Bowker & Star who describe boundary objects as “those objects that both inhabit several communities of practice and satisfy the informational requirements of each of them” (1999: 297). Such an object appears robust enough to travel across contexts and between communities of practice in an identifiable form, yet flexible or “plastic” enough to take on the

meaning of the local context. Different groups can and do inscribe different meanings on the information represented in a specific artefact or process. Thus, a boundary object can serve a key role in developing and maintaining coherence across communities. Further, the boundary object as an analytical device draws attention to the possible form such “objects” might take. This conception of boundary objects has received wide acceptance in a number of disciplines (see Anderson 2007 for further discussion). For Roth & McGinn (1998) boundary objects are inscriptions used across communities of practice and constitute embodied representations. They serve as ‘interfaces between multiple social worlds and facilitate the flow of resources (information, concepts, skills, materials) among multiple social actors’ (Roth & McGinn, 1998: 42). The focus turns from representation as mental activity to inscription as social activity. We build on these conceptions to examine ways artefacts, processes and players in the project served such functions within and across the communities involved.

Research Project as Boundary Object

What we have in Fair Tracing is an attempt at maintaining a logical but awkward space of spaces. There is value in keeping openness in many dimensions because of the type of tool being proposed, yet when these are considered cumulatively; it is to produce a potential miasma. In particular we can identify:

- the openness of meaning necessary for Development work across business processes, cultural boundaries and understandings of knowledge to be valuable to any participant;
- the openness of approach necessary to ensure that everyone from different research and practice communities can contribute ideas and perspectives;
- the openness of participation boundaries to ensure that all related chain actors can contribute their defining perspective and seek their own representation;
- the openness of the design solution so that all the learning from the project may be finally bound into one unifying system or many distributed ones;
- the openness of the technology so that the knowledge/code can be adapted.

Clearly, this creates a highly indeterminate problem space. This would be a weakness on a development project, but this was a research project. We can regard this indeterminacy as an opportunity. To make this point, we have selected three contrasting perspectives on the project to review for their implications.

The Los Robles version of the Fair Tracing idea

The Fair Tracing partner in Chile was the Los Robles winery collective⁵ and their suppliers and owners, 44 vineyards of whom four are Fair Trade certified. Producers in Chile serve goods to supermarkets in Europe and America that have stringent accountability demands. Because they are part of these supply chains as well as others, Chilean producers have to meet global standards for exported goods, even if they are not required to meet them domestically. Producers are

⁵ Unfortunately, the collective was dissolved in 2008 because of economic pressures.

dependent on major distributors for their livelihood and it becomes in everyone's interests to be able to isolate any batch quickly and eliminate any problematic items (and *only* those items that are problematic or might be contaminated).

The winery has a precise audit trail. The logistics manager writes in log files that trace each bottle of wine back to a date of production, a vat of fermentation and a batch of grapes. The log files exist in large bound books that reside behind him in his office and that can be cross-referenced in a matter of minutes if there is an issue with any of the bodega's output. Upstairs the oenologist is blending wines and recording her results in a dedicated database called Kupai, which she shares with the reception centre that grades producers' output into A, B and C category grapes. In the next room, the lab staff are measuring acidity and putting the chemical analysis into a Word document that is then printed and stored in a folder. In other words, the use of ICT for mapping the supply chain within the bodega itself is fractured and involves multiple types of record. Only some are in a form that can be manipulated. The story is much the same for the growers. Some have spreadsheets for recording data; some, a book. One foreman records information in an Excel spreadsheet for himself, then takes the data out of it to interact with the rest of the vineyard's production processes.

Although the bodega employees were interested in the idea of a marketing tool for communicating with consumers and prepared to consider what kind of information they might enter into it (Fig 1B), they talked about tracking technologies. They saw a means of putting data into a Fair Tracing tool from each stage of production and in so doing bypassing the fragmentation of their internal systems. For them, the real value of the tool was for logistics management.

The Indian coffee sector's version of the Fair Tracing idea

The Indian coffee sector is not of one mind. Industry bodies such as the Indian Coffee Board, a national government organisation, and CoMark, a marketing cooperative of coffee growers from the three coffee-producing Indian states, each have an agenda. Speciality coffee growers, and the federations and many self help groups that support very small growers (with land of 10hectares or less) have their place. Economic circumstances such as falling coffee prices, poor harvests over 10 years, deregulation of the industry and the tax situation (if you sell over the gate of the estate you pay no tax, but if you take produce to the curing works you pay 25%) determine growers' willingness to take some actions and not others.

Two trends are relevant here. Indian organisations are keen to use technology, in keeping with India's mission to be a country at the forefront of science and technology. They are also aware of a trend to go it alone, without international (and particularly global multinational) intervention.

Following coffee from the small plantations to the curing works, as noted above, is to watch it lose its identity immediately. Whereas Los Robles wine is accountable to the last drop, much coffee makes its way from plantation to sack to curing works – often already in the hands of an independent small trader – and from there into multi-source instant coffee powder (roasted and packaged, but not in such a way that a thread unites a producer with the jars on supermarket

shelves). Small paper chits record the amount of different grades of coffee so that growers can be remunerated. Mechanisation of data is very limited. To all intents and purposes, the Indian coffee sector did not look ready for the tool, with no records to grab automatically as part of constructing a chain and no traceability.

The impact of the Fair Tracing project in India was unexpected then. Without much interest in the research project, the different interests in the sector came together to work on tracing. United by the efforts of one PhD student to understand and affect the conditions of coffee production, as of December 2008 a working party involving people up to state government was investigating tracing technologies for the whole industry. There was no intention (at time of writing) of widening the remit to other sectors beyond coffee, extending its scope beyond India, or of involving the research team in the development of the tool.

The Ethical Consumer Information System (ECIS) version of Fair Tracing

In late 2008, researcher M2 hosted a research seminar on “Ethical Consumption, Traceability & ICT”. Many members of the team were there, joined by others interested in ethical consumption. Presentations during the day involved several not-for-profits, Web2.0 outfits and a small commercial company active in the tracing space. A major topic was that no organisation had solved the problem of representing ethical information in a form that indicated at a glance what was being meant by ‘ethical’. Each system seemed to have met only part of the challenge. Some had user-generated recommender modules; some had scoring for different ethical conditions. The Fair Tracing project brought ideas of provenance in the form of working with producers and representing the value chain.

At time of writing, thirteen British, European and American organisations with an interest in promoting ethical consumption were meeting virtually in a bi-weekly phone conference to discuss data standards. Some of these groups had been at the seminar, but others became interested through word-of-mouth over the following weeks. Many participating organisations are not interested in research as such, but are pooling knowledge and ideas about a consumer tool and common data/information standards. Drawing on these speculative discussions, participants with the try-it-and-see philosophy of Web2.0-cum-agile-methods make a working prototype between meetings. Approaches are evaluated and modified. Most of these projects are fuelled by individuals’ private passion and it is symptomatic that the commercial company could not afford to take part. The outcome will not be Fair Tracing research; it will be the joint discoveries of the ECIS.

Analysis: the project as player

It is clear from the stories above that the ‘bridging’ function of the Fair Tracing idea – particularly in the shape of the value chain – proved a powerful device around which local narratives could form. The chain metaphor worked in two ways across contexts: to make explicit the production process as both a physical and virtual phenomenon and to stress the connection to trade and consumption activities. We might say that the value chain functioned as a ‘meaning probe’.

Other features, such as the rhetoric of “Open Source” were also persuasive, and reflect the moral commitment of Development. In short, the idea proved a tractable representation of a series of categories, or system, with the correct moral gloss to evoke interest and win alignments – a catalyst. The focus on inscription offered by the notion of the boundary object helps us appreciate this meaning-making as social activity moving across and between these communities.

Other Boundary Objects

Looking at Fair Tracing is like staring at a fractal picture. With the degree of openness yet definition apparent in this project, it is possible to show how almost all aspects were boundary objects of one form or another, but here we pick two more to analyse in detail. The first, the value chain, is a major part of the project as both process and artefact, while the second, the blog, is a minor component but shares the symptomatic characteristics of the whole.

Value Chain as artefact and process

Returning to the value chain, we now explore its function more closely. It is possible to see this production and transportation process through the purposes of players in the project. One of the main participatory exercises conducted (for instance, during a Chilean design workshop, Fig 1B) was asking partners to draw a version of the value chain and annotate it with information they collect and would consider of interest to pass on to consumers. The following hopes for computer-supported collaboration appeared during discussions:

- To represent the value chain as a means of educating the consumer
- To represent the chain as a means of educating the players in the chain
- To connect the end-producer to the consumer (and vice versa) for dialogue
- To connect the production actors in the chain to each other
- To give (easier) accountability to producers in dealings with distributors
- To give specific social, environmental and economic information on products
- To identify *individual instances* of items and track them across their journey
- To alter power relations between actors in the chain
- To eliminate parts of the chain which add no value but take resources.




Fig 1A: the interaction designer’s interpretation of the stories told her about the chain. Fig 1B: Bodega employees in Chile draw out the value chain and annotate it for a consumer audience.

We might note that, while no two items are mutually exclusive, their range is impressive. Even designing an underdetermined artefact would make the final tool potentially very complicated. And taking the value chain as a design stimulus didn't simplify matters for the team, as their multiple perspectives reveal:

From computer science: the chain is a series of sequential events that can be bounded, identified and ordered in terms of time and duration between transitions. Transitions occur at each point that value is added, ownership changes or movement initiated. Each product and actor can be given a discrete number that will allow products to be mapped and assembled in chains. Information can then be attached to any event in a variety of multimedia forms through the creation of fields. (design meeting, UCL, London)

From economic geography: "The term value chain describes the full range of activities which are required to bring a product or service from conception, through the intermediary phases of production to delivery to final consumers, and disposal after use (Kaplinsky 2001). Different value chains are characterised by processes and actors involved, by the nature of linkages between these actors, and by the overall governance structure of the chain (Gereffi et al 2005) [...] including the related power relations, information flows and discourses". (Kleine 2008:110)

From interaction design/HCI: a series of physical journeys between stages in the production process involving the use of multiple technologies and different information systems, which more or less efficiently ensure the translation and perpetuation of information as the product passes towards the consumer. At each stage a different cultural, economic and political context determines business people's tolerance for additional data entry overheads, the kind of interfaces that are appropriate and the form of representation information should take. (design meeting, home of interaction designer)

From IT entrepreneurship: friends and contacts in a mesh of social relations that make up Bangalore society, who embrace certain values, will band together if interests are best served, who have astute relations with members of the state government and the national trade board for their commodity, who share information strategically, and who will consider making their product traceable if benefits outweigh tax penalties. (fieldwork, Bangalore)

Analysis

The project deliberately set out to collect alternative readings of the value chain and these fed into the design of the consumer work to investigate spatial, temporal or social representations. A principal objective was to collect information about what producers were prepared to share in terms of confidentiality of business data, collection overheads, self image, etc. The team also sought to hear partners express their relation to the rest of the chain, noting the political issues. We can see a range of political motivations in the list, from educating others to improving local knowledge to eliminating actors who siphon off profit. However, the exercise of constructing the chain was also means to get partner teams talking together and negotiating their understanding of the issues. In this respect, the artefact (the value chain idea) inspired a new process (constructing the chain together) which mirrored the actual process of progressing along the chain. This

process involved identification and articulation of chain functions, politicizing them. And this chain became artefact again through drawing (Fig 1A and B), to become both a record and a further tool with which to negotiate meaning.

It is also clear from the accounts above that no single understanding of the chain existed in the research team (though the accounts are polarised versions). It is possible to see a mix of applied and abstracted description. The CS version is stark; others are progressively more social. Some versions of the chain presuppose an intervention while others reflect a greater interest in how chains exist and run through society. The difference between design- and social science-orientated disciplines reveals itself. But, in reality, although everyone adopted different positions as a starting point, individuals gave accounts drawing on different traditions at different times for different purposes. Particularly the HCI specialist (M3), who bridged subgroups, used multiple positions to keep a broad view of the issues alive. And moments of definition helped clarify problems.

For instance, in reviewing the possibility for an end-to-end numbering system based on events (the CS view), it became more apparent to the team that not only was the overhead of data entry too much to ask of small businesses (certain, but not all, data could be automatically 'grabbed'), but there would be problems fitting third parties in. Bodies like the Fair Trade Labelling Organisation which certifies Fair Trade goods or Oxfam couldn't operate in the chain using numbers, codes and events and would probably have to be 'fed in' through trade bodies or producers as intermediaries to attach them to the right parts of the right chain. The choice between creating a system to be used by single producers to communicate directly with consumers, or one that could only operate as a federalist syndicate reflecting a whole sector came into relief (Dearden and Light 2008). So, engagement with the engineering view threw attention on the socio-economic.

The blog

The other example to be examined here is less central to the research work of the project, but had a critical place in raising profile, building credibility and disseminating. The Fair Tracing blog (www.fairtracing.org) served as a main repository for any information or ideas in the team that were seen to have lasting significance. Further, the blog was public-facing, and comments, messages to team members and personal encounters reveal that it was read by interested people without a direct link to the project. The resource had major responsibility for giving the project an external image. But the blog was a compound object, free of syntheses or summaries. Many postings concerned tracing technologies and organisations interested in ethical production or sourcing, but others recounted Fair Tracing activities or those of Development projects using e-commerce. Encyclopaedic style entries sat next to stories, while other forms of writing, such as the abstracts of academic papers, and media, such as some video and extensive images, accompanied more narrative sections. Postings were tagged with themes that related to the purpose of the project, but new tags could be initiated by anyone – more folksanomic than hierarchical – and the chronological presentation of the blog meant that a pot-pourri of information met the casual visitor. In summary, the many voices of the project owners joined up as one

aggregated vision of making change, but did so as a collage, not a set of intentions or goals.

Analysis

The Fair Tracing blog, as a small part of the bigger whole, reflected the multiplicity of versions of the project in a typically Web 2.0 way, where plurality is enshrined and folksonomies challenge the ordered world of information management and hierarchies. By eschewing a single authorial position, it gave purchase for all comers to the project through its diversity and lack of single interpretation, but was still distinct enough to be part of the branding of the project as recognised by the team members. In a sense it was a perfect avatar of the project: pluralistic in nature and never committing to one perspective. It offered an interesting contrast between the coherence of individual posts published (in styles acknowledging the different writing traditions of different members' fields as well as preferences) and the arbitrary composite of posts viewed. With no overall design, extracted meaning is entirely situated and emergent. Its plurality encourages the multi-interpretation possible of the project, shown above in other respects.

Interpretation and Reflection: Boundary as Bridge

Much analysis has already been included. Here we look at the overall positioning of the project and what this might contribute to further work in the field, given our intention to support understanding of working with ICT in a Development context and to extract value from an 'amorphous' interdisciplinary research project.

The development process reflected an aim to work sensitively across cultures – reflected it, perhaps, too closely. Prototyping went on but without precise recommendations for an end-to-end tool. We have noted the breadth of the research question (entertaining much production, trade and consumption worldwide), the limit of researchers' time and reach (a factor of financing rather than personal competence, but nevertheless relevant), their ambitions to work collaboratively with their partners and to keep the bigger 'bridging' vision of the project in sight. Against this backdrop, individual findings, such as acceptability of methods and the reception of interfaces (presented in Light et al 2009, Light 2008, Kleine 2008, but not addressed here), become almost insignificant.

In asking about the difference between two production contexts, the team noted that a technical solution could not be simple. In asking about the needs and desires of different partners, the team established that the purpose of a Fair Tracing tool could be as diverse as the actors it might represent. In asking about the messages that were seen as important to communicate along the chain, the team recognised that any tool would never be a fixed representation but would forever be a site of contested meaning, made complex by subtleties of language, values, forms of expression - challenges brought to the fore in cross-cultural

work. In seeking to represent a socio-technical system to that system, the team accepted that there was no chance to be dispassionate observers; that the act of asking questions was as influential as any action research and that, before any actual implementation had an impact on relations, the fact of the research would enter and change the system.

In sum, the output of the research, when judged against other UK research projects, was deemed successful using various multi-disciplinary success criteria, including negatives like not exploiting representatives of small business. But was there value in trying to keep the bigger 'bridging' vision of the project in sight, when tackling any smaller part might have been more productive in conventional design terms? What kind of research team holds tenaciously to an idea even when they find that the pursuit of it stymies greater creativity at the implementation level? And is there anything from the experience that might feed into other projects to enable them to function more effectively, given that indeterminacy and multiplicity of interpretation sit at the heart of Development work?

The team collected a wealth of inconsistent but useful design information. The diversity of it speaks to creating a system with the openness seen in the project blog, described above, where voices from different actors form patterns in the eye of the beholder. However, openness is only useful if it is sufficiently structured that it can be exploited. For instance, the ECIS example shows the value placed on consistent standards in ethical data to underpin user rankings. And structuring proved contentious, since different actors had different purposes in wanting the tool and thus appreciated different functions. In our initial examination of issues, we pointed to values: deciding what is best for people's wellbeing and how this decision-making is managed between players (funders, researchers or the political hierarchy of intended beneficiaries). To ask whose wellbeing should be targeted is to ask whose functionality to embrace. In doing Development work, change for the greater good is a higher order purpose than embedding tools and thus not for an exogenous team to legislate upon, any more than they should presume to know *a priori* what is important locally. We have discussed some knots into which this agenda can tie a project; and we have pointed to Fischer's idea of developing environments (not solutions) to provide an interpretive background (2007). What we argue now is that the negotiations of the project offer a possible – just one possible – navigation of the political and ethical aspects of this wilderness.

Here we return to Erickson's question: How do designers begin when they are unsure of what they are making, what it should do, or who will use it? (1995). The ethical path would not seem to be the decisive one. It would be fair to say that the project boasted an ill-defined problem: a design landscape one forms as one goes through it (Dorst 2003). But it might be appropriate also to describe it as 'ill-defined research' in the sense that, at start, it is not clear *who* is the subject of the research, or more accurately, who it is appropriate to sample, and it is not clear *what* the research is to provide. The team began a discovery phase with participatory elements and never truly emerged from it, despite conducting valid research. The team eschewed an early synthesis of stakeholder perspectives and consequently never synthesised them. But they also created the space of spaces for political discussion and appropriation to take place.

There were ever more versions of an idea, each with its own constraints, and behind it always the same vision. Using the device of the boundary object we have shown how appropriation of the vision became possible. This would be of gentle theoretical interest were we not able to point at two startling appropriations (India, ECIS) that could not have been predicted. While one cannot deliberately design boundary objects, in using the boundary object as conceptual device for analysing the project and its many elements, we are reminded of the role that inscription of meaning plays in trans-disciplinary and trans-cultural contexts like those experienced throughout the activities described in this paper. Thinking in terms of boundary objects and boundaries gives one cause to ask which boundaries are being crossed and how they are maintained. Sensitivity to conditions and relations, the intention to include partners in defining and articulating the research question and desire to deliver something of moral value all position the research, despite its all-too-typically exogenous origins. It is, then, not so much that there is something to learn from the project, as recognising how tenacity and flexibility permeated it and to what effect. It is offered as a project to think with when designing other interventions. Presented here, in yet another interpretation – i.e. as we have seen it and accompanied by our analysis – it draws attention to what each set of actors perceives to set them apart (unique values) and, in this way, suggests a responsive approach to issues that are not unique to Development. Indeed, we expect to see similar issues pervade a greater number of projects if, as promised, pervasive computing brings digital technology to colonise more aspects of life.

In summary, through this analytical process, we are made mindful that artefacts are socio-material forms. Similarly, examples from the Fair Tracing project also show how process can be framed as such socio-material objects. For this reason we have come to recognise that the final product, so to speak, of the project was, in fact, a process that has been capable of holding form across contexts and communities while remaining plastic enough to take on the values and meanings inscribed by local communities and players. The device of the boundary object helped us realise that, while this condition naturally poses many challenges for the ‘design’ goals that motivated the project, the result is a necessary consequence of holding steadfast to the core values of designing collaboratively in Development – values that remained constant across all contexts.

It is interesting to reflect, and this paper begins to do so, what happens in a design research project for a socio-technical system that seeks to occupy multiple spaces with enough interpretive flexibility to allow it to become meaningful and embedded in the wider world. The strength of the idea (the connection between producer and user, the value chain, the drive for ethical behaviour, the moral tool) has qualities of the boundary object. And the way that different team members formed a loose aggregate provided the open weave that allowed the boundary object to endure in an endlessly pre-designed – and therefore fertile – condition.

Acknowledgments

The authors thank the Engineering and Physical Sciences Research Council, and the “Bridging the Global Digital Divide” network (esp. Andy Dearden for his comments), and acknowledge the work of Fair Tracing’s Ian Brown, Ashima Chopra, Dorothea Kleine, Apurba Kundu, Helen Le Voi, Cote Montero, Macarena Vivent and Christian Wallenta, and timely provocation by Paul Dourish.

References

- Anderson, T.D. (2007): ‘Settings, arenas and boundary objects: socio-material framings of information practices’, *Information Research*, vol. 12, no. 4, paper colis10. Available at <http://InformationR.net/ir/12-4/colis/colis10.html>.
- Bentley R. and Dourish P. (1995): ‘Medium vs. mechanism: Supporting collaboration through customisation’, in H. Marmolin, Y. Sundblad and K. Schmidt (eds.): *ECSCW’95*, Kluwer, Dordrecht, pp 133-149.
- Bowker, G. C., & Star, S. L. S. (1999): *Sorting things out: classification and its consequences*, MIT Press, Cambridge, MA, USA.
- Dearden, A. and Light, A. (2008): ‘Designing for e-Social Action: An Application Taxonomy’, in Proceedings of the Design Research Society Conference 2008. Available at: <http://www3.shu.ac.uk/Conferences/DRS/Proceedings/Proceedings.htm>.
- Dearden, A. and Rizvi, H. (2008): *Adapting Participatory and Agile Software Methods to Participatory Rural Development, PDC’08*, Indiana University Press, Bloomington IN, pp. 221-225.
- Denzin, N.K. (1989): *Interpretive interactionism*, Sage, Newbury Park, CA, USA.
- Donaldson, K. (2008): ‘Why to be Wary of “Design for Developing Countries”’, *Ambidextrous*, Spring 2008, pp. 35-37.
- Dorst, K (2003): *Understanding Design*, BIS, Netherlands.
- Dunne, A and Raby, F. (1999): *Hertzian Tales*, Royal College of Art, London, UK.
- Erickson, T. (1995): ‘Notes on design practice: stories and prototypes as catalysts for communication’, in J. Carroll (ed.): *Scenario-based design: envisioning work and technology in system development*, John Wiley & Sons, Inc., New York, NY, pp. 37-58.
- Fischer, G. (2003): ‘Meta-Design: Beyond User-Centered and Participatory Design’, in J. Jacko and C. Stephanidis (eds.): *HCI’ 03*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 88-92.
- Fischer, G. (2007): ‘Designing Socio-Technical Environments in Support of Meta-Design and Social Creativity’, *CSCW’07*, Rutgers University, July, pp. 1-10.
- Gaver, W.W., Boucher, A., Pennington, S. and Walker, B. (2004): ‘Cultural probes and the value of uncertainty’, *Interactions*, vol.11, no.5, pp 53-56.
- Greenbaum, J and Kyng, M (eds.) (1991): *Design at work: Cooperative Design of Computer Systems*, Lawrence Erlbaum Associates, Hillsdale, NJ.
- Holden, J. (2004): *Capturing Cultural Value*. Demos, London. Available at: <http://www.demos.co.uk/files/CapturingCulturalValue.pdf>.
- Irani, L., Dourish, P., Grinter, R. and Phillips, K. (working paper, 2009): *Postcolonial Computing*.
- Kleine, D. (2008): ‘Doing action research on Chilean Fairtrade wine value chains’, *Geographical Journal*, vol. 174, no. 2, pp. 109–123.

- Light, A. (2008): 'The Challenge of Representing a Sociotechnical System: Fair Tracing and the Value Chain', *Sociotechnical Aspects of Interaction Design*, London, May 2008.
- Light, A. and Luckin, R. (2008): 'Social Justice and User-centred Design', *Futurelab Opening Education series*.
- Light, A., Kleine, D and Vivent, M. (2009): 'Performing Charlotte: a Tool to bridge Cultures in Participatory Design' *Int. Journal of Sociotechnology and Knowledge Development*, 1(2)
- Lincoln, Y.S. and Guba, E.G. (1985): *Naturalistic inquiry*, Sage, Newbury Park, CA.
- Muller, M.J. (2002): 'Participatory Design: the third space in HCI', in J. Jacko and A. Sears (eds.): *The Human-computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, Lawrence Erlbaum, Hillsdale, NJ, pp1051 – 1068.
- Nardi, B.A. & O'Day, V.L. (1999): *Information ecologies: using technology with heart*, MIT Press, Cambridge, MA.
- Porter, M.E. (1985): *Competitive Advantage: Creating and Sustaining Superior Performance*, Free Press, NY.
- Rittel, H. and Webber, M. (1984): 'Dilemmas in a General Theory of Planning', in N. Cross (ed.): *Developments in Design Methodology*, Wiley & Sons, Chichester, pp. 135-144.
- Roth, W.-M., and McGinn, M. K. (1998): 'Inscriptions: towards a theory of representing as social practice', *Review of Educational Research*, vol. 68, no. 1, pp 35-59.
- Star, S. L. (1996): 'Working together: Symbolic interactionism, activity theory, and information systems', in Y. Engeström & D. Middleton (eds.): *Cognition and communication at work*, Cambridge University Press, Cambridge, UK, pp. 296-318.
- Star, S. L. and Griesemer, J. R. (1989): 'Institutional ecology, 'translations' and boundary objects: amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39', *Social Studies of Science*, vol. 19, no. 3, pp 387-420.
- Sterling, B. (2005): *Shaping Things*, MIT Press, Cambridge, MA, USA.
- Suchman, L. (1987): *Plans and situated actions: the problem of human-machine communication*, Cambridge University Press, Cambridge, UK.

Supporting Nurses' Information Flow by Integrating Paper and Digital Charting

Charlotte Tang and Sheelagh Carpendale
University of Calgary, Canada
char.tang@ucalgary.ca; sheelagh@ucalgary.ca

Abstract. Information technology has changed the way health care is delivered. Electronic health records which are prevalently deployed to replace or supplement paper documentations have made distributed information access at various points of care and work activity achievable with the use of mobile information devices. Our particular concern is with nurse's information flow, where nurse's notes and observations taken at the point of care feed into the electronic record. In these cases, digital technology has not yet entirely replaced paper and pen, because the latter still provide greater ease and flexibility of use when compared to current digital technologies. Even when mobile digital technology is available, clinicians still prefer creating handwritten notes, and then later manually transposing them into the digital medium. Within this context, we created a prototype that integrated digital paper with electronic health charts to retain the benefits of paper and pen, as well as digital medium. A focus group evaluation of this prototype demonstrated promise and potential for its value in a medical environment.

Introduction

Our work concerns nurse's *information flow*, where there is discord between the way nurses use pen and paper to record notes and observations taken at the point of care, vs. how they feed this information into the electronic record. The focus on the flow of information from a nurse into the Electronic Health Record (EHR) is a crucial factor in overall nurses' effective collaboration and coordination. If the nurse can enter primary notes and readings during patient contact, the nurse/patient face to face interaction is naturally facilitated (i.e., note taking will not interfere). If the information is entered into the EHR in a correct, complete and timely way, the way other nurses on the current or on a later shift can exploit

this information (usually asynchronously) is also improved. Of course, the long-term archival record is itself an information source that is used by many others, be it administrators, doctors, and so on. Again correctness, completeness and timeliness is critical.

Currently, many clinicians persistently rely on paper personal artefacts (e.g., personal notes written on a note pad and carried around) that they informally use in the course of their work (Hardey et al. 2000, Fitzpatrick 2004, Tang & Carpendale 2006). This is despite advances in information technology in the last few decades that have considerably changed the way health care is delivered. The importance of paper personal artefacts to support effective information flow in patient care has been well recognized (Allen 1998, Hardey et al. 2000, Silva et al. 2006). Notably, they offer flexibility of use with ongoing tasks (Luff et al. 1992, Mackay 1999, Sellen & Harper 2002, Nomura et al. 2006). Newly emerging information can be easily added to these paper artefacts during work in progress. In contrast, the process of updating information via digital devices such as PDAs or Tablet PCs is slower (Silva et al. 2006). Studies also found that nurses relied heavily on their paper personal artefacts when delivering patient care, even when alternative mobile information technology was available, e.g., a computer-on-wheels with bedside information access (Tang & Carpendale 2008). In fact, many nurses in our study ward regarded these paper personal artefacts as indispensable to their nursing work and non-replaceable.

At some point, this information must be fed into the EHRs. Currently, EHRs have increasingly replaced or supplemented paper-based documentation. These digital records are important. They allow quick information access across distributed locations without the need for physical transportation of paper information artefacts, which was often time-consuming, required careful planning, and needed coordination among multiple parties. With the increasing awareness of the importance of having information easily available at points of care for improved patient care, various mobile technological devices have often been introduced, e.g., PDAs, smart phones, Tablet PCs, and computers-on-wheels (Silva et al. 2006, Tang & Carpendale 2006, Zamarripa et al. 2007). However, these devices rarely afford an intuitive mode of interaction for accessing information at points of care (Cohen & McGee 2004, Lu et al. 2005). Thus, paper personal artefacts are still heavily in use. Nurses often create these at the beginning of a shift. They extract information from distributed information sources creating a copy on paper. For instance, information source could be computer terminals, individual patient charts, or verbal communication from colleagues (Tang & Carpendale 2006, 2008).

The mandatory process where nurses document medical information that emerges during nurse's shift in formal records is called *charting*. Currently, charting is time-consuming and constitutes a considerable segment (around 13-28%) of nursing activities (Allen 1998, Strople & Ottani 2006). As part of their

information flow, nurses have to manually transpose the information from their handwritten notes on paper artefact into the digital EHR. This is not only time-consuming but is also prone to errors. (Zamarripa et al. 2007).

Given this context, we set ourselves the research goal to design technological solution to support nurses' actual work practices (Tang & Carpendale 2006, Vicente 2004) around their information flow. Our prototype allows nurses to maintain their familiar and efficient practice of personalized information recording on a paper-like interface, while integrating their paper inscriptions and direct input into the digital hospital information system. We believe such paper and digital integrated charting has the potential to improve not only charting efficiency, but will increase information timeliness by making that information immediately available to other clinicians. We conducted a focus group study around our prototype as a preliminary step to evaluate its effectiveness in facilitating nurses' information flow practices.

Our technology design is based on the findings from in-depth field studies in a local hospital ward (Tang & Carpendale 2006, 2008) combined with iterative design discussions with the nurses. To set the scene, we introduce our design process by first briefly recounting the findings of these two studies, paying particular attention to shift changes during which the paper personal artefact is created. These findings are gathered to formulate goals for designing applicable technologies, leading to a description of our integrated charting concept. We then describe the focus group study we conducted to evaluate our technology prototype. This is followed by our findings. Finally, we summarize the lessons learned as adjusted and refined design guidelines for building technologies to support the nurses' information flow practices.

The contributions of this research are:

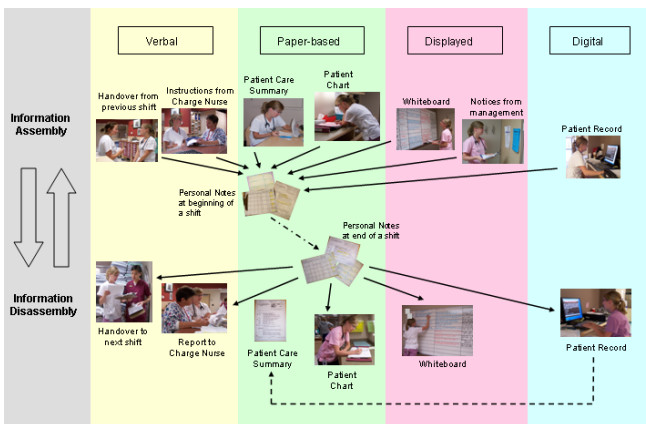


Fig. 1: Incoming nurses acquired information from distributed multimedia information sources (top row), placed the extracted information in their paper personal notes (centre), added new information to personal notes, and finally charted the new information (bottom row).

- The identification of the need to explore and improve the interrelationships between nurses' informal information practices and the organizational information system in a hospital setting.
- The design and prototyping of a paper and digital integrated charting concept, designed to support informal use of paper personal artefacts as an integral part of the official documentation process.
- The study and analysis of focus group feedback from nurses to evaluate the effectiveness of this integrated charting approach.
- The presentation of a detailed set of design guidelines for future technology development.

Design Process

In this section we describe how our field studies (Tang & Carpendale 2006, 2008), and design discussions with the nurses working on the study ward led to our design ideas and the creation of paper and digital integrated charting.

Baseline Observational Study. The first study provided us with a thorough understanding of how information flow took place during nurses' shift change and through their shift, as summarized in Fig. 1 (Tang & Carpendale 2006). Nurses coming on a new shift had to acquire a good understanding of the current operation in order to carry on the patient care appropriately. To facilitate this, nurses used an intermediary artefact, typically a note sheet which we simply refer to as personal notes (Fig. 1, centre), which has been found to play a crucial role in nursing work. They wrote on their personal notes a consistent set of information that they extracted from distributed multimedia information sources such as patient charts and the EHRs. This graphological practice has been repeatedly reported to help build a mental model of their nursing duties for the shift (Kidd 1994, Tang & Carpendale 2006). Surprisingly, this manual exercise to record information at shift change was not included in the handoff strategies recommended by Patterson et al. (2004). Nurses also customized their personal notes with individualized layout of information that helped them to retrieve information from their written markings (Fig. 2, left). In the course of their shift work, nurses added newly emerging information to their personal notes and used it to constantly adjust their work plan. Finally, the gathered information had to be properly documented to support continual patient care across consecutive shifts (Fig. 1, bottom).

In brief, nurses relied upon these personal notes for specific functions: as their work plan for nursing work, as their immediate information source and opportune notepad during their shift, and as their information basis for reporting and handover for uninterrupted patient care (Tang & Carpendale 2008). Thus it is important to design technology to preserve these benefits and to support individual nurses' needs for personalized note taking.

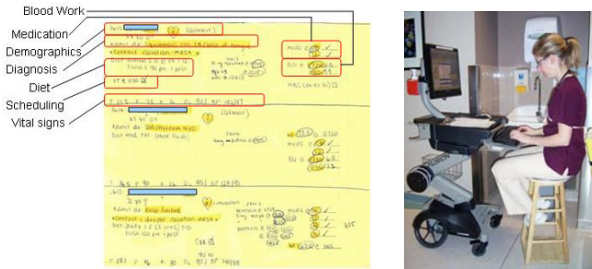


Fig. 2: (Left) Personal notes with customized layout. (Right) A computer-on-wheels.

Studying the Impacts of a Mobile Technology. In the second study we investigated how the deployment of a mobile information technology (a computer-on-wheels as in Fig. 2, right) impacted the nurses' information flow (Tang & Carpendale 2008). The mobile technology introduced in our study ward was often found not to be used for information access at points of care as intended. Contrasting with the affordances offered by paper notes, these mobile technologies fell short in supporting the important roles necessary for carrying out nursing functions as described above. This motivated us to focus on how the paper and the digital world might be bridged so that information can be converted efficiently across the divide while preserving the benefits of paper use.

Paper and Digital Integrated Charting

Findings from the field studies were validated by our participants, through proofreading papers and by providing feedback at presentations. Our design ideas evolved iteratively in response to feedback received from the participants through discussions and presentations. To illustrate iterative nature of the design process, the baseline study identified that extracting information from the EHR and writing personal notes on paper was time-consuming and tedious. Thus, we first intended to replace this manual transcription practice with a computer interface which would allow nurses to drag and drop required information from the EHR directly onto a personal notes template and provide features to augment important information such as highlighting or using different colors. These personal notes can be printed, used as immediate information source, and as an opportune notepad adding information as it is acquired during their shift. However, subsequent informal discussion with the nurses indicated that the manual handwriting process at the beginning of a shift was crucial for building a solid mental model of the agenda for their shift (Ash et al. 2004). Similarly, feedback was continuously fed into our design process leading to re-design, re-thinking, or validation of different aspects of our design. Our design goals are:

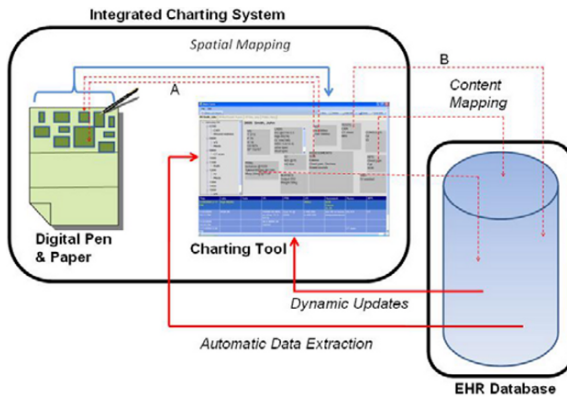


Fig. 3: Design scheme of the integrated charting system.

- To build upon the affordance of paper,
- To support nurses' current information flow practices,
- To use the advantages of EHR to provide instant availability of consistent information to distributed clinicians

Paper and Digital Integrated Charting Overview

In our *paper and digital integrated charting* (PDIC) approach, the EHR exists in full functionality, uncompromised and, if desired or if necessitated by circumstance can be used and interacted with as before. In addition, information gathered from the EHR can be organized with digital pen and paper as best suits personal work practices. This digital paper personal information artefact can be used through the shift just as previous personal paper notes were used, recording work done and collecting pertinent information as it arises. A charting interface facilitates the recording of the information generated during the shift into EHR.

A schematic design of the integrated system as integrated with EHR is shown in Fig. 3. It adds two parts to the EHR information system: *digital pen and paper technology* (Fig. 3, left & Fig. 4), to allow nurses to use tangible pen and paper as they did previously, and a computer application, *Charting Tool* (Fig. 3, center & Fig. 6), for creating digital paper templates and converting handwritten notes into




Fig. 4: (Left) Personal notes template printed in digital paper with customized layout of bounding boxes and static information (Right) A digital pen docking station.

digital text (Fig. 6, top). The dotted lines (A & B) in Fig. 3 indicate the mappings between these components and are described below.

Paper and Pen Digital Integrated Charting Components

Digital Paper and Pen

Our PDIC makes use of the Anoto™ digital pen and paper technology. Regular paper is printed with Anoto dot patterns which when combined with a pen that contains a camera provides digital functionality. The patterns consist of numerous small black dots that are digitally legible and form an innocuous background that does not interfere with normal use for writing and diagrams. The existence of many patterns can make individual pages uniquely identifiable. The digital pen (Fig. 4, left) contains a digital camera that captures all the markings made on digital paper. The dot pattern on each digital paper indicates the exact position of markings made by a digital pen on the paper. This location-sensitive capability of the digital paper supports creation of distinct regions that can be identified visually for organizing information and computationally when interfacing with the EHR. It is also possible to discriminate between information written in a region or between regions. Thus all notes including information on the exact position the markings on the paper are digitized by and stored in the pen. This information can be downloaded to a computer immediately via Bluetooth technology, or at any point in time using the pen docking station (Fig. 4, right).

Creating personal notes templates. Individual nurses can customize digital paper for their shift by using the *charting interface* (see Fig. 6) to create and save a template that specifies the information types and their layout. To create a template, an information type, e.g., vital signs, is chosen and a box that bounds the region that will contain this information type appears in the charting interface. It can then be moved and resized as desired. This method provides flexibility for nurses to each has their own type of personal notes and binds the regions on the digital papers with the specified information type. Other information types can be added in the same way until the template layout is satisfactory. Thus the charting interface can be used for creating any number of information layouts. Based on findings that nurses used a layout for all patients and preferred all their patients' information displayed on a single page (Tang & Carpendale 2006), once a template is created nurses can specify the number of patients and obtain a digital paper that uses their template sized appropriately for their number of patients. For example, Fig. 5 shows a template created on the charting interface (left), and two personal notes for 4 (middle) and 5 (right) patients using the same template for different number of patients. This reflects the way nurses dynamically adjust their personal notes in practice.

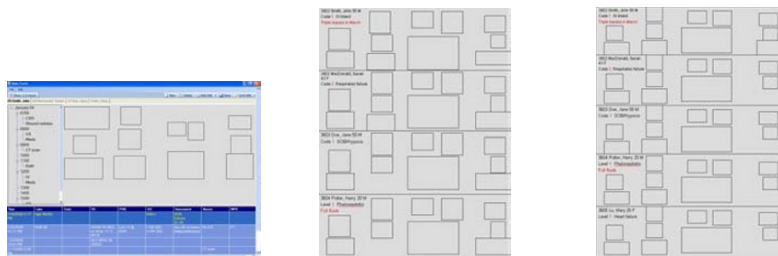


Fig. 5: (Left) Personal notes template created on Charting Interface (Middle & Right) Personal notes printed for 4 and 5 patients respectively

Using digital-paper personal notes. At the beginning of a shift, nurses print their personal notes based on their template on a digital paper (Fig. 5, middle & right). Patients' static information such as demographics and medical history can be printed with the template (Fig. 4, left A). Nurses can then use their digital personal notes with their digital pen. If desired they can still highlight important notes with regular highlighters (Fig. 4, left), as has been common practice. Thus these digital personal notes can be used as before to facilitate nurses' work both as a bed-side information look-up, and for recording newly emerging information. For example, as before nurses can record vital signs in the vital signs region.

Nurses' normal practice also includes notes used for references or for personal use that is not intended to be charted. For this, PDIC provides two alternatives: 1) spaces outside the bounding boxes are not charted by default, or 2) bounding boxes can be designated for this specific purpose. Handwritten notes placed in region thus specified will be discarded and will not be converted into digital text for charting, whereas handwriting in the other bounding boxes will be processed and converted into digital characters for charting.

Charting Tool

The Charting Tool is the software interface between the digital paper and pen and the EHR. It consists of three components (Fig. 6):

- a single interface *charting interface* for editing and charting (top)
- a *quick reference* to view up-to-date clinical information (bottom)
- a *task timeline* to support planning of the shift work (left)

To allow a nurse to focus on one patient at a time, the interface displays information for one specific patient, switching between patients is easy with the tab design (top row).

Charting Interface. The same charting interface is used by nurses to create personalized digital paper notes and to transfer handwritten notes captured with the digital pen to a computer. Connection between the digital pen and EHR is done either via the pen docking station or continually via Bluetooth during nurses' shift. However, to maintain the integrity of the EHR the downloaded information will not be processed until the charting interface is used to verify it. The downloaded notes are converted by handwriting recognition software, into

digital characters which are displayed in bounding boxes corresponding to the defined regions on the digital paper. In the charting interface these notes can be edited if necessary. Nurses are required to verify and confirm the correctness of the information to be charted before saving into the EHR. Once confirmed, the charting interface bounding boxes are mapped to the corresponding fields in the EHR database (Figure 3, B). Through this charting interface, nurses do not have to navigate through the EHR and instead of manually entering the data for charting they simply verify and edit, potentially saving a lot of nurses' time. The charting interface also allows nurses to chart spontaneous but important information from sources other than their personal notes. This can be added to appropriate bounding boxes or entered into a new clear box. Information entered in this manner also requires verification.

Provision for resizing and relocating the bounding boxes, and for dragging and dropping text between bounding boxes allows nurses to easily move information to the correct bounding box. The interface is also customizable and extensible for other less-frequently used information types in order to meet the needs of the patients. For example, if blood work is not already in the template, nurses can dynamically add a bounding box labelled *blood work*. Nurses can then type in information pertinent to blood work, and this will be saved to the corresponding field in the EHR along with other information.

Quick Reference. The quick reference provides an at-a-glance overview of the EHR as pertains to the patient in view with most recent information first. This enhances information retrieval from the EHR and can also be used to countercheck if specific data has been charted. New information added to the EHR via the charting interface will be instantly displayed and highlighted in the top row of the quick reference. This quick reference can also be customized in



Figure 6: The charting tool consists of a charting interface, a task timeline, and a quick reference to medical records.

that rows can be highlighted if desired and the columns displayed can be customized for specific types of information.

Task Timeline. On the left hand side of the charting tool is a timeline of tasks to be performed for the currently selected patient that is initially based on time-based information such as medication at certain time intervals that can be automatically extracted from the EHR. This timeline provides an optional planning tool for clinicians to organize their shift work, can be used as a quick overview of shift work, and as a reminder and to-do list. Nurses can also dynamically adjust the timeline (i.e., add, edit, remove, move tasks) in response to changes in their patients' needs and their own temporal horizon in the progress of their shift work (Reddy et al. 2006). In addition, it can also aid reporting to e.g., charge nurse or an incoming colleague, as it displays the tasks that were planned and/or achieved.

The Study

To evaluate *paper and digital integrated charting*, we prototyped the basic concepts and ran focus groups with participants from our original study ward. Focus groups were chosen to provide a good initial step in iterating the system design as participants could acquire hands-on experience with the prototype and freely express and discuss their experience and expectations.

Participants

Focus group participants were nine registered nurses and three nursing students. They all were experienced in preparing and using paper personal notes for their shift work. Because the participants were the targeted end-users of the technology, they were highly motivated to offer their practical experience and expectations of the technology. Six focus groups were conducted in total. Data from one group was not used because participants opted out of key portions of the focus group protocol. The participation in the focus groups was negotiated around existing nurses' schedules. These groups were planned for times when two or more people agreed to participate. The final composition of the focus groups ranged from 1 to 5 participants (5, 3, 2, 1, 1) due to unexpected illnesses and family emergencies.

Method and Materials

Each focus group lasted about an hour and participants engaged in these activities.

1. Creating an individual template. Each participant drew on paper the spatial layout of their own paper personal notes, labelling the respective types of

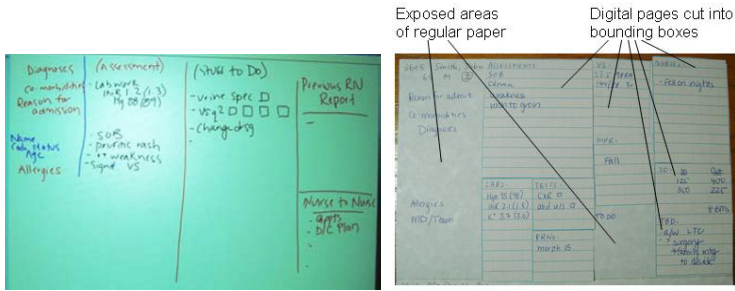


Figure 7: (Left) Collaborative template created on Smart Board. (Right) Corresponding paper prototype created by gluing snippets of (lined) digital paper on a piece of regular paper (plain grey). information such as medications and assessments. Stationery such as pens, markers and highlighters of different colors were provided. Participants were encouraged to create the template as close as possible to the personal notes they created and used in practice. This activity took less than five minutes.

2. Creating a collaborative template. Participants worked as a group to discuss and design a template for personal notes that was acceptable to all members of the group using either a Smart board or a large flip-chart (Fig 7, left). This group template was needed for use in the next activity. Single participants simply used their own template. This collaborative template can also help inform the design of a generic template for use by nurses who have not yet created their own template, e.g., new nurses.

3. Prototyping personal notes. Pages from Adapx™ Digital Journalx, which were lined and printed with Anoto™ dot pattern (Fig. 7, right), were used. These pages have the same property that PDIC is based on. Each page is identifiable due to its dot pattern. Participants then cut pages to create required regions and glued them on a piece of regular paper (Fig. 7, right) to match the layout of the collaborative template from activity 2 (Fig. 7, left). Thus notes and handwriting inside each region created from digital-paper will be captured, digitized and identified for a specific kind of information by a digital pen. Exposed areas of the regular paper (Fig. 7, right) could be used for information not required to be charted.

Minor inconsistencies were found in the information types predefined in the charting interface and those used in the collaborative template created in Activities 2 and 3. For example, separate regions were provided for laboratory work and tests whereas some collaborative templates contained only a combined “labs and tests” information type. Therefore, participants had to slightly adjust the layout of the bounding boxes on the constructed digital paper notes. These constructed digital paper notes were in turn used in the next activity.

4. Exploring the PDIC prototype. The researcher explained how the boxes and areas outside the boxes on constructed digital paper notes interfaced with the charting tool. The charting tool was displayed on a Smart board mounted to the wall (Fig. 8) to offer a clear view to all the participants and to trigger discussions

incorporating their own practical experience and expectation (Wilson et al. 2006). Participants were frequently encouraged to discuss and comment on their experiences and expectations of the technology to better support their work. How the digital pen and paper work was also explained such as the dot pattern imprinted on the digital paper, the embedded camera in the digital pen, and why the exact position of markings made on digital paper was known.



Fig. 8: Participants exploring the integrated charting.

Participants then used a digital pen to write on their paper note (Fig. 7, right), pretending that they were preparing and using it as their personal notes in a work shift. They were encouraged to write as they would normally do. Information sheets which contained fabricated medical information, previously compiled with the help of a registered nurse, were provided for participants to prepare the digital personal notes. However, all participants wrote down information from their experience without consulting the information sheets.

The handwritten data was then downloaded via the pen docking station (Fig. 4, right) connected to a computer. The researcher triggered the use of handwriting analysis software to convert the handwritten data into digital text and explained how this happened (i.e. Wizard of Oz technique) whereas this would be automatically done in a fully-implemented system. Participants could then click the 'Notes' button on the charting tool and see the text converted from their handwriting that was inside the bounding boxes on their digital paper notes. Participants were encouraged to edit and add more detailed information to the displayed text to the point that it was appropriate for charting. They then saved the information after verifying its correctness on the verification screen. Participants were encouraged to explore and comment on all aspects of the charting tool such as the timeline and the quick reference update.

Findings

The main goal of the study was to identify strengths and weaknesses of the design of paper and digital integrated charting, in particular whether it can facilitate or impede nurses' work practices around their information flow, and to gather suggestions for improving the system design. The study employed a combination of varying-fidelity methods instead of a full-fledged system to demonstrate the design concepts and to actively draw the participants into the design dialog. The

first three activities, i.e. to draw a personal template, to sketch a collaborative template and to craft a paper prototype of personal notes using digital paper snippets, were conducted to progressively prepare for the last activity which was the focus of the study. Therefore, we briefly highlight the findings from the preparatory activities and concentrate on the findings of the last activity.

Highlights of Preparatory Activities

Create an individual template. Although the individual templates appeared different, they all contained a consistent set of information: static information (room number, patient name, age, medical history, doctor's name, resuscitation level, allergy, special diet, diagnosis), and dynamic information (vital signs, assessments, I/Os, IV, tests, labs, blood work, medication, previous shift report, and new orders). All our participants confirmed that they created and used paper personal notes for their work.

P1: "... I would never give up my piece of paper. Literally if I lose it, I'd spend like 40 minutes trying to find it. I'll go through every garbage bin 'cause it's like your lifeline. I don't know what to do with myself if I don't have it."

P4: "Because our entire day works on that piece of paper, everything we do, when we need reports, we wouldn't go to the computer. It's all on that [the personal notes]."

P6: "I don't know what to do whenever I lose it [the personal notes]."

Create a Collaborative template. Three focus groups conducted this collaborative activity. Two groups created a new collaborative template based on discussion of the participants. One group selected one of the participants' personal templates as their group template after comparing and contrasting their personal templates and finding that the spatial layout was the only difference among them. Moreover, many participants said that they made changes to their personal templates in terms of both the layout and the information content from time to time when their experience increased and when their patients' needs differed. The activities and discussions around designing a collaborative template indicated that a default but adjustable template could be useful.

Prototyping personal notes. All groups constructed their digital paper notes using snippets of digital paper glued on a regular paper. Three groups conducted this activity based on their collaborative template and the other two were based on the participant's personal template.

Exploring the Paper and Digital Integrated Charting Prototype

Participants' responses towards our technology design were mostly positive. They were particularly delighted that the design was based on pen and paper interaction. They saw great potential and value of the system in facilitating their work. Several participants asked when the technology would be available for use

in practice. Yet, they also identified weaknesses of the system and made suggestions for improving the technology design.

Strengths

Pen and paper. Participants' reliance on their paper personal notes during shift work was in evidence as they referred to personal notes as their 'lifeline' and 'my entire day works on [my personal notes]'. They perceived great benefits in retaining the use of the familiar pen and paper and were excited to see that their handwriting could be converted to digital text without requiring manual re-typing at a computer. They voiced hope that this could improve their charting efficiency as they would only need to edit or add to the converted information.

Single-interface charting. Participants found the ability to chart multiple information types with the same interface useful. They also found the flexibility to dynamically extend the interface to include other kinds of information for charting convenient. They expected that charting would become more efficient and more comprehensive because current tedious and time-consuming navigation of the hierarchical EHR could be minimized. Making current information pertaining to one patient visible on the same interface provided an overview and a more comprehensive picture of patient's condition that could help more easily organize and document different kinds of information gathered during the shift. They also commented about huge time savings to benefit both actual nursing care and hospital finance by paying less overtime work.

'Personal' information space. Some participants wrote everything inside the digital-paper bounding boxes on the constructed digital personal notes whereas others used the space on the regular paper for information that was not meant for charting (Fig. 7, right). The former pointed out the need for a 'personal' information space on their personal notes when they saw the converted text of their personal notes displayed in the charting interface included also their 'personal' information. They explained that this kind of information was vital for accomplishing their nursing work, regardless if it is information extracted from documented sources to serve as reference information, or their own personal opinions towards a patient, e.g., as a reminder of specific ways to deal with the patient, or just some scribbles for catharsis to relieve their stress. In short, all the participants agreed that personal notes should provide 'personal' information spaces for writing down information that will not be charted.

Focus on new information. Participants liked the design that previously charted information is not displayed in the charting interface so that they can always focus on newly emerged information. This is because the charted information would still be readily available in the quick reference table below. One participant suggested including an option for displaying the charted information in the charting interface as long as it was easy to differentiate between the old and new information such as using different color or different font style.

Quick reference facilitates information retrieval. All the participants regarded the quick reference that offered an at-a-glance view of the archived medical information useful. They found it convenient to look up medical records without having to go through the hierarchical EHR. They also found the dynamic updates in the quick reference table when new information was saved to the EHR helpful in that they could be more aware of up-to-date patient information.

P1: *"I like it [quick reference] because... sometimes you don't get a good report from the previous nurse, then you can scroll back to see what the nights [nurse in night shift] think. I wish we had something like this... For example, if a patient fell, the nights tells the days and the days tells the evenings, but when nobody tells anybody anything, you forget that this person may fall. But if there's this thing [quick reference], I can just go back to see what happened and like what's the plan 3 days ago."*

P12: *"I really like that I don't have to go to different places to look for new things for the patients and it's all here in the blue table [the quick reference table]. I like that!"*

Timeline useful to new nurses. Most experienced nurses did not find the task timeline useful for planning and organizing their shift work.

P11 (experienced nurse): *"nursing uses flexible process... constantly shifting... no way to plan like this... If you're going to spend time on this, you'll never get to patient care".*

However, student nurses liked the timeline. They perceived it a useful tool to help them focus on their work and to allow them visualize their work schedule and work load. Many experienced nurses recalled that they also manually created similar timelines for planning their work when they were new. But they had abandoned this practice as they became more experienced with their nursing duties and were more used to the dynamic and flexible nursing processes (Bjørn & Balka 2007). Therefore, they believed that the timeline would be a valuable tool for training and educating novice nurses.

Weaknesses and Concerns

Liability issues. Participants were particularly concerned about their liability for the technology in case they lost their digital pen. Since a digital pen is much more costly than a regular pen which nurses do not worry about losing, participants were worried if they had to be financially responsible for replacement. This led to active discussions on ways to prevent losing the digital pen. Suggestions included carrying it around their neck with a lanyard but they were then worried about the number of gadgets that they had to carry with them. They were also concerned about the technology's costs, durability and environmental issues (e.g., recyclability of digital paper).

Handwriting recognition. Our participants did not express negativity towards the handwriting conversion because their handwriting was generally converted quite well in the study. However as expected, many participants were concerned with the general quality of handwriting recognition. They were also worried about their handwritten symbols and abbreviations (e.g., SOB for short of breath and CXR for chest x-ray) and whether they could be properly transcribed.

Embedded charting. Participants criticized the system in that it did not facilitate ‘flow sheet charting’ which is an embedded structure in the EHR. They are required to access this by navigating the information system in order to chart by “point and click” a checklist of information categories in the flow sheet. They suggested having a button that linked with the flow sheet so that they could directly access it without navigating the information system. The researcher proposed using digital-paper flow sheet printed with checkboxes so that charting on flow sheets can be done anywhere and the checked information can be easily transferred to the EHR. This alternative was well received, especially by experienced nurses who prefer paper artefacts to digital medium. But younger nurses found a single-mouse-click link to the digital flow sheet more beneficial.

Suggestions

Extending the potential of digital pen and paper technology. In view of the potential of digital pen and paper, participants suggested designing the personal notes to include areas for other documentation and reporting purposes in addition to EHR charting. As currently, nurses have to write or verbally report the same information multiple times in different places such as in patient charts, verbal shift reports and large whiteboards, requiring much mobility and redundancy of effort. Therefore, they considered it a great value if they only needed to write once on their personal notes and the information would be saved or displayed (with or without converting to digital text) instantly in different media. For example, participants could write a shift report on their personal notes for relevant clinicians such as their incoming colleague or charge nurse to retrieve at a computer. Participants perceived significant savings in time and mobility when they no longer needed to look for people in order to give or receive reports.

Support both customizable and customized interfaces. Many participants resented that they always had to navigate to their last visited screen every time they logged on the EHR. They were also frustrated that the current EHR did not allow them to keep their customized views of information which would be automatically reverted to the default view when they logged off the system. They estimated that this practice of logging on and off the EHR took place 50 to 100 times per shift. Thus they emphasized it is imperative that our system is designed to allow them to continually use their customized views and to display their last visited screen upon logging on. They also expected that these features would greatly improve their work efficiency.

Discussion

Integrating the effective aspects of current work practices with the advantages of the EHR is the goal of our Paper and Digital Integrated Charting. Thus this research commenced with observational studies to form a thorough understanding of current work practices. The evidence from our field studies indicated that while

nurses' use of pen and paper held many advantages, it was in conflict with the hospitals' goal of the consistency of the EHR. This led us to design technology using *digital pen and paper* so that handwritten notes can be easily digitized, thus working towards the goal of integrating the best of both directions.

The focus group study we conducted using a prototype developed with mixed-method approach provided encouraging feedback. For example, "*This is really quite exciting and I know there're glitches that are not working as well as we'd like. But we're moving towards the right direction*" (P11). While weaknesses were identified and concerns were expressed, the feedback gathered indicated potential of PDIC and gathered together provides the following set of refined guidelines.

Support flexibility and personalization. Participants in general were most pleased by the potential for both supporting their current work needs and having the flexibility to support changing work practices.

- The use of flexible templates, as exemplified by the use of bounding boxes to create identifiable and recognizable interface components that can be resized and reorganized to match personal preferences provides both consistency for the computational interface and individuality for personal use.
- Familiarity with interfaces can facilitate work (Kidd 1994). Thus it is important that the customized interfaces can be saved for ongoing and future use.

Facilitate information entry and retrieval. Participants discussed several aspects about current systems that were time-consuming from disseminating information to multiple media to repeated clicking in the EHR. They were hopeful that PDIC would provide considerable time savings.

- Providing navigation that is based on recognizable interface components which have established system mappings has considerable potential for streamlining information updating and retrieval.
- When possible embed required information structures (e.g., flow sheet, multidisciplinary report) within an integrated interface.

Provide an overview. Participants liked being able to see the information gathered from their digital personal notes in conjunction with the timeline and quick reference. They particularly appreciated the quick reference and its instant update on information verification and entry.

- Combining several information representations in a customizable interface avoids problems with fragmented information that has been scattered over different places and has been found to impede the building of mental model (Ash et al. 2004).
- Viewing a comprehensive set of information can provide feedback on what has been done and still needs to be done thus greatly facilitating the collaborative effort across shifts for providing quality patient care, as well as the charting process and its quality.

- Enriching the information presentation, such as including the timeline which was found to help novice nurses plan, organize and focus on their work, may further enhance the learning outcome and the work efficiency.

Support ‘personal’ information use. Personal notes, as our participants indicated, are created and used by their owner so are inherently personal and traditionally contain both archival and non-archival information that is only intended for personal use. However, with the use of digital pen and paper, the real ‘personal’ information would run the risk of being publicized over the digital medium.

- Some information spaces should be assigned strictly for ‘personal’ use only without leaving any digital trace. Otherwise, the technology would fail to support individual needs and may experience adoption resistance.

Safeguard the accuracy of information. This factor is always a concern for health care systems and our nurse participants are no exceptions. Many discussions and comments focused on this topic primarily in three ways.

- The system must always instigate a mandatory, yet lightweight, verification safeguard before information is saved to the database to uphold the integrity of information accessible by distributed clinicians.
- The system should provide an option of keeping a cache of personal notes information that can be retrieved by its owner. This will be useful in case the frontline information artefact needs to be reproduced.
- The reliability of handwriting recognition is a justifiable concern. However, the use of mandatory verification makes this feasible and advances in handwriting recognition technology, the captured, digitized handwriting can be converted into digital text with increasing accuracy, narrowing the divide between paper and digital medium.

Support system dispersal. Our participants noted, and is confirmed in previous studies, that they both retrieve and disseminate information from/to multiple sources. These rich varieties of information artefacts in medical settings facilitate various groups of clinicians to accomplish many different goals (Bardram & Bossen 2005). Thus it is not uncommon that the same information has to be inputted in multiple media. This practice undoubtedly cost time and effort.

- An integrated charting system that is linked with other information artefacts and displays could minimize redundancy of effort (Cabitza et al. 2005) and to provide information in multiple modals and representations (Reddy et al. 2001). This in turn will benefit patient care.

Conclusion and Future Work

Based on previous studies on nurses’ information flow, we have designed, prototyped and studied a paper and digital integrated charting solution that offers:

- manual transposing of information from multimedia sources to a digital paper notes for building a mental model and planning of the shift work,
- the ability for individual nurses to customize their digital personal notes for facilitating information retrieval,
- portable, flexible and low-cost use of digital paper notes to support bedside information access and note-taking,
- easy transfer of information from digital paper notes to the EHR to provide timely, low-cost and continuous information flow,
- the use of personal notes as information basis for reporting in non-digital media such as in verbal shift reports,
- a quick reference to archived medical information, and
- a timeline for visually planning work.

Through focus groups of practising nurses, we studied this technology and obtained valuable feedback on the benefits they perceived and well-articulated suggestions for improving the system. Together, they helped to refine our set of design guidelines that other researchers and designers may find helpful in their specific settings. In order to discover how to best support nurses' work practices, we have taken an iterative design, prototyping and evaluation approach. The next step of this research would be to re-design and re-implement the integrated charting system based on the refined guidelines.

Acknowledgement

This paper has greatly benefitted from the valuable input of Saul Greenberg. We also thank the nursing staff at W21C, Foothills Hospital, Calgary, and our funding agencies, AIF, iCore, NERSC, and SMART Technologies for their support.

References

- Adapx. <http://www.adapx.com>, on Jan 20, 2009.
- Allen, D. (1998). Record-keeping and routine nursing practice: The view from the wards. *Journal of Advanced Nursing* 27:1223-30.
- Anoto Technology. <http://www.anoto.com> on Jan 20, 2009.
- Ash, J. S., Berg, M. and Coiera, E. (2004). Some Unintended Consequences of Information Technology in Health Care: The Nature of Patient Care Information System-related Errors. *Journal of the American Medical Informatics Association* 2004; 11 (2): 104-112.
- Bardram, J. and Bossen, C. (2005). A Web of Coordinative Artifacts: Collaborative Work at a Hospital Ward. *Proc. of GROUP*, pp. 168-176.
- Bardram, J. and Bossen, C. (2005). Mobility Work: The Spatial Dimension of Collaboration at a Hospital. *JCSCW* 14(2), pp. 131-160.
- Bjorn, P. and E. Balka (2007). Health Care Categories have Politics too: Unpacking the Managerial Agendas of Electronic Triage Systems. *Proc. of ECSCW*, pp. 371-390.
- Cabitza, F., Sarini, M., Simone, C. and Telaro, M. (2005). When Once Is Not Enough: The Role of Redundancy in a Hospital Ward Setting. *Proc. of GROUP*, pp. 158-167.

- Cohen, P. and McGee, D. (2004). Tangible Multimodal Interfaces for Safety-Critical Applications. *Communications of the ACM* Jan 2004, Vol. 47 No. 1, pp. 41-46.
- Fitzpatrick, G. (2004). Integrated Care and the Working Record. *Health Informatics Journal* 10(4), pp. 291-302.
- Hardey, M., Payne, S. and Coleman, P. (2000). 'Scraps': hidden nursing information and its influence on the delivery of care. *Journal of Advanced Nursing*, 2000, 32(1), pp.208-214.
- Kidd, A. (1994). The Marks are on the Knowledge Worker. *Proc. of CHI 1994*, pp. 186-191.
- Lu, Y., Xiao, Y., Sears, A., Jacko, J. (2005). A Review and a Framework of Handheld Computer Adoption in Healthcare. *International Journal of Med Informatics* 74(5), pp. 409-422.
- Luff, P., Heath, C., and Greatbatch, D. (1992). Tasks-in-Interaction: Paper and Screen Based Documentation in Collaborative Activity. *Proc. of CSCW 1992*, pp. 163-170.
- Mackay, W. (1999). Is paper Safer? The role of Paper flight Strips in Air Traffic Control. *ACM Transaction on computer-Human Interaction* 1999, 6(4), pp. 311-340.
- Nomura, S., Hutchins, E. and Holder, B. (2006). The Uses of Paper in Commercial Airline Flight Operations. *Proc. of CSCW*, pp. 249-258.
- Patterson, E., Roth, E., Woods, D., Chow, R. and Orlando, J. (2004). Handoff strategies in settings with high consequences for failure: lessons for health care operations. *International Journal of Quality in Health Care*, 16(2), 2004, pp. 1-8.
- Reddy, M., Dourish, P. and Pratt, W. (2001). Coordinating Heterogeneous Work: Information and Representation in Medical Care. In *Proc. of ECSCW*, pp. 239-258.
- Reddy, M., Dourish, P. and Pratt, W. (2006). Temporality in Medical Work: Time also Matters. *JCSCW* 15(1), 2006.
- Sellen, A. and Harper, R. (2002). *The Myth of the Paperless Office*. The MIT Press, Cambridge.
- Silva, J., Zamarripa, M., Strayer, P., Favela, J. and Gonzalez, V. (2006). Empirical Evaluation of a Mobile Application for Assisting Physicians in Creating Medical Notes. *Proc. of the 12th Americas Conference on Information Systems*.
- Strople, B. and Ottani, P. (2006). Can Technology Improve Intershift Report? What the Research Reveals. *Journal of Professional Nursing*, 22(3), 2006, pp. 197-204.
- Tang, C. and Carpendale, S. (2007). An Observational Study on Information Flow during Nurses' Shift Change. *Proc. of CHI*, pp. 219-228.
- Tang, C. & Carpendale, S. (2008). Evaluating the Deployment of a Mobile Technology in a Hospital Ward. *Proc. of CSCW*, pp. 205-214.
- Vincente, K. (2004). *The Human Factor – Revolutionizing the Way we Live with Technology*. Vintage Canada Edition 2004.
- Wilson, S., Galliers, J. and Fone, J. (2006). Not All Sharing is Equal: The Impact of a Large Display on Small Group Collaborative Work. *Proc. of CSCW*, pp. 25-28.
- Zamarripa, M., Gonzalez, V. and Favela, J. (2007). The Augmented Patient Charts: Seamless Integration of Physical and Digital Artifacts for Hospital Work. C. Stephanidis (Ed.): *Universal Access in HCI, Part III, HCII 2007, LNCS 4556*, pp. 1006-1015.

The eDiary: Bridging home and hospital through healthcare technology

Rikke Aarhus¹, Stinne Aaløkke Ballegaard² & Thomas Riisgaard Hansen¹
Centre for Pervasive Healthcare, Department of Computer Science¹, Department
of Information Studies², Aarhus University
raa@cs.au.dk, imvsab@hum.au.dk, thomasr@cs.au.dk

Abstract. The main contribution of the paper is to present challenges relating to the use of new healthcare technology, the eDiary, which seeks to create a better integration between home and hospital. To minimise risks of malformations and other complications, pregnant women with diabetes are enrolled in an extensive treatment regime, which requires frequent visits to an outpatient clinic as well as a high degree of self-care. The eDiary is designed to assist the women in this work, primarily by allowing the women to register their glucose values, record video consultations, and support video-tele-consultations. This paper reports on a pilot study during which pregnant women with diabetes and their healthcare providers make use of the eDiary. The pilot study indicates that such healthcare technology not only allows the women to achieve a better integration of the management of their diabetes into their everyday life, but may also challenge existing power relations between patients and healthcare providers.

Introduction

The demand for hospital services is increasing as new treatments, lifestyle related diseases and a growing elderly population require more interventions. To move treatment from hospitals to the patients' home and to empower the patients are considered possible responses to this challenge. However, moving hospital services into the homes and everyday lives of patients has a number of consequences for patients in how they manage their disease. It also has consequences for the healthcare providers because it changes organisation of the

healthcare services, collaboration between different healthcare providers, and the collaboration between patients and healthcare providers (Dinesen 2007).

The focus of this paper, and the project HealthyHome it is based on, is to enquire into challenges related to the implementation of new healthcare technology that integrates home and hospital. HealthyHome was a two-year, Danish research project focusing on the design of healthcare technology to be integrated in the everyday life of people living at home with a health condition. A secondary focus was to bridge health-related activities in the homes with activities at the hospital. The project was a joint project between a university, two industrial partners and a hospital. One of the industrial partners specialized in electronic health records and the other in wireless technology. The case studied in the project was pregnant women with diabetes. They matched our objective of working with patients that required extensive health support from the hospital, but still spent most of the time out of the hospital. The participating healthcare providers and secretaries were from the outpatient clinic where pregnant women with diabetes are treated.

Based on the development, implementation and pilot study of a web-based tool, the paper points to aspects related to shift in workload, responsibilities and power relations between the home and the hospital. First, we will briefly introduce the case, the developed tool and the pilot study, and then discuss the findings from the pilot study.

Pregnant women with diabetes

A pregnancy is a complicated condition for women with diabetes as diabetes causes risks of pre-eclampsia, hypertension, premature birth, malformation of the heart, miscarriages, and stillbirth (Lauenborg et al 2003, Jensen et al 2004, Clausen et al 2005). According to one of the obstetricians involved in the project, approximately 50 % of the pregnant women with diabetes have a Caesarean birth as their foetuses weigh too much for a normal delivery.

To reduce the risk of complications, the pregnant women with diabetes are closely observed by a specialised interdisciplinary team of healthcare providers (dietician, diabetes doctor, obstetrician, and midwife) during pregnancy and birth. Since 2001, the treatment has been centralised at four specialised units in Denmark (Indenrigs og Sundhedsministeriet 2003). The treatment consists of consultations with the team of specialists at one of the four units every second week until week 32 and then every week¹. In addition to the hospital-based treatment extensive self-care is required to keep the blood glucose low. Due to the pregnancy the need for insulin fluctuates making it difficult to keep blood glucose stable. The self-care includes exercise and a healthy diet.

¹ A normal pregnancy is expected to last 40 weeks, however labour in diabetic pregnancies is often induced during week 37.

In the beginning of the project, a field study was carried out focusing on this double-sided treatment of diabetic pregnancies. The study lasted four months and included approximately one hundred hours of observations at the outpatient clinic as well as interviews. Ten pregnant women with diabetes were recruited at the clinic and interviewed in their home. Four healthcare providers were interviewed at the hospital. The ten pregnant women all had type 1 diabetes, as opposed to the type 2 diabetes lifestyle related, and the majority of the women have had diabetes for several years. The field study revealed the main challenges within the existing treatment programme and of being a pregnant woman with diabetes. In particular the amount of work required by the women in carrying out their self-care and attending the consultations at the outpatient clinic informed the further design process.

Related work

Easy and correct management of blood glucose is generally a huge challenge for people with diabetes and hence a large number of both commercial systems and research projects address this area. This section will briefly discuss some of the software solutions and systems that address everyday management of diabetes.

A comprehensive list of a couple of hundred freeware and commercially available systems are listed and briefly commented by David Mendosa (Mendosa, web 2009). Most systems provide similar functionality and are mainly focused on visualising glucose level, insulin doses, calorie intake and exercise information. A challenge for the use of these systems is to enter the required information into the system. To address this a number of the solutions provide extensive information about different types of foods (DiabetesPilot, web 2009), support mobile data entry (SiDiary, Mendosa, OneTouch, web 2009), or allow synchronization of data from for instance a glucose meter (AccuCheck, Onetouch, web 2009). However, even though the data entry is easier with some of these solutions, they still require extensive data entry activities from the user, which reconfigure some of the solutions from being about decision support and overview to a registration tool as discussed by Danholt (2008).

A number of the commercial systems also use the title diary or logbook to emphasise the continuous use of these systems, but still they focus strongly on health or diabetes and do not mix everyday events with the clinical purpose of collecting data (SiDiary, MyNetDiary, DiabetesLogBook, web 2009).

A number of research projects have also worked on the management of diabetes. A larger endeavour is the MAHI research project by Mamykina et al (Mamykina, 2006, 2008). The MAHI project extends a previous project called CHAP within this area and focuses on people newly diagnosed with diabetes. In the MAHI project 25 people used a combination of a camera-phone and a glucose meter for four weeks to record glucose levels and take pictures relevant to their diabetes. The main focus of the project was to help people reflect on how to

manage their diabetes. It is a different challenge from the pregnant women with diabetes where most have had diabetes for several years.

A related project investigates the relation between digital photos and glucose data. In this project a system with a glucose meter and a camera was tested in a pilot study (Smith, 2007). And while the focus is on creating a tool to support reflection, the project shares the same challenge as the previous project. To really make sense of the collected data, the data needs to be coupled to the context. For instance the interpretation of a blood glucose figure depends on whether the data is taken before or after lunch. And though photographing events such as eating lunch might help the interpretation this might not be a viable solution for the everyday use of the software to manage the blood glucose.

Design, implementation and evaluation of the eDiary

In line with Mamykina et al (2008), we believe that rethinking health records is relevant with the rise in chronic diseases where being ill is an aspect of everyday life and with the increasing number of treatments being moved from hospital to home entailing collaboration between patient and healthcare provider outside the hospital. Management of disease and treatment becomes a matter of supporting the integration between hospital and home and facilitating the treatment at home. In this section, we present the process of design as well as the eDiary developed to achieve just this.

The field study showed that pregnant women with diabetes generally experienced that the management and treatment of their disease took up much space and time and were, to some extent, what their everyday life evolved around during their pregnancy (Ballegaard & Aarhus 2009). The purpose of the eDiary is hence to be a tool for supporting the pregnant women with diabetes to manage their diabetes in their everyday life, but also a tool for supporting their collaboration with their healthcare providers in a hospital setting, as this was another point from the field study. The eDiary mixes the concept of a diary and a personal health record to achieve this objective. A diary is often a personal item for recording everyday events relevant to the owner. In contrast a personal health record is often the outcome of clinical systems of medical devices and contains everyday information about the patient's health condition.

The concept and the following prototype implementation were developed on basis of the early field studies and on a series of workshops with participation of project partners, healthcare providers and pregnant women with diabetes. Based on the early field studies, ten different concepts were presented in an initial design workshop with healthcare providers, two pregnant women with diabetes and the research team. Among the concepts were: specially designed handbags for pregnant women with diabetes, an intelligent booking system, and novel

consultation rooms. Based on the discussions at the workshop, the concept of the eDiary was selected and further developed.

To validate our discussions, a second workshop was designed to get feedback of the eDiary from a larger number of patients. The workshop took place in the waiting area at the outpatient clinic. A third and fourth workshop explored a mock-up of the eDiary and tested a prototype of the system. Four pregnant women with diabetes and their healthcare providers (an obstetrician, two midwives, a diabetes doctor, a dietician and two nurses) participated in these workshops.

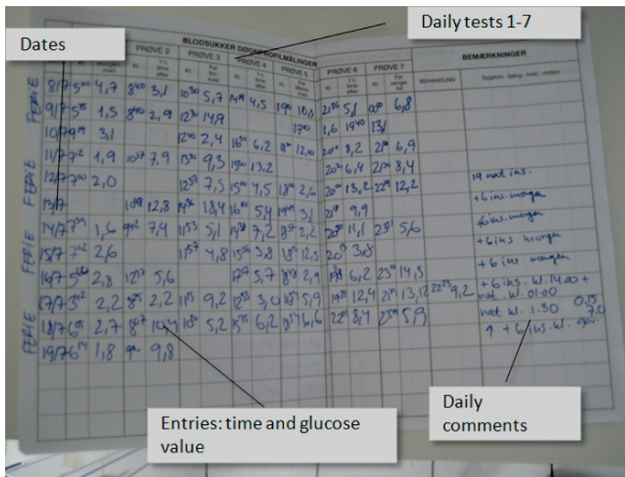


Figure 1. A traditional diabetes book. Each row presents the blood glucose development during a single day, along with insulin dose and comments.

The eDiary prototype

The eDiary consists of a web portal, which offers various services. First, the portal allows users to manually type in blood glucose measurements from their preferred glucose meter and add additional comments to specific measurements. The field study revealed that comments were vital for understanding the circumstances of a specific measurement – if eating birthday cake caused a high figure or if it indicated a shift in the need for insulin (see (Dourish 2004) for a discussion of the use of contextual data in system design). Furthermore, the portal contains a monthly overview of blood glucose measurements similar to the traditional diabetes book that the pregnant women receive from their doctor shown in Figure 1. A traditional diabetes book. Each row presents the blood glucose development during a single day, along with insulin dose and

comments. In this book they register their blood glucose values approximately seven times a day, note their insulin dose, and add comments, if necessary. They use the book in their daily evaluation of their blood glucose and in the discussion with the diabetes doctor at the outpatient clinic. Apart from this, the portal can automatically highlight with colours women, e.g. all figures below 3.5 or higher than 10.² Some blood glucose meters allow automatic upload of data, however, this approach was not selected because there was no easy way of attaching comments to the individual readings and because the women used devices of different brands, many of them not able to connect to a computer.

Screenshot 1

eDagbog

Dates Columns: Before breakfast | 1 ½ hour after | before lunch ...

DATO	Før morgenmad		1½ time efter		Før frokost		1½ time efter		Før aftenmad		1½ time efter		Før sengelid		Om natten		Kommentarer								
	KL	BG	Insulin	KL	BG	Insulin	KL	BG	Insulin	KL	BG	Insulin	KL	BG	Insulin	KL		BG							
1/10	06:30	16,1		09:00	8,2		12:00	6,9		14:30	6,8		17:30	8,3		20:00	9,8		21:00	14,8				Syg 10 ^a -2-7-2-11 ^a	
2/10	06:00	16,2		09:00	5,4		12:00	6,7		17:30	8,7					20:00	4,7		21:30	9,1					10 ^a -8-7-8-11 ^a
3/10	07:00	8,0		09:00	3,9		12:30	7,4		14:30	6,2		17:30	5,1		20:00	4,7		21:30	6,8					10 ^a -8-7-7-11 ^a
4/10	8:00	8,2					12:00	4,9		14:00	4,4		17:30	3,3		20:00	3,1		21:30	16,6					10 ^a -8-7-7-(12 ^a)
5/10	04:00	4,2					12:00	9,7		14:30	6,2		18:00	16,4		20:00	11,2		21:30	8,6					10 ^a -8-8-9-12 ^a
6/10	07:00	3,3		09:00	9,2		11:20	10,2		14:30	3,3		17:30	6,9		20:00	5,0		21:00	4,8					2:00-4,9. 10 ^a -7-8-8-2-12 ^a
7/10	07:00	6,1		09:00	8,2		12:00	9,2		16:00	9,1		18:30	13,6		20:00	9,4								10 ^a -8-8-8-12 ^a
8/10	06:30	6,4		09:30	3,3		11:30	9,1		14:00	3,6		17:30	5,7		20:00	5,4								10 ^a -8-8-2-8
9/10	06:00	8,3		09:30	17,4		12:30	2,7		14:30	2,4		17:30	9,4		19:30	2,4		21:30	6,8					10 ^a -8-8-8-12 ^a
10/10	06:30	6,4		09:30	6,7		12:30	5,8		14:30	2,8		18:00	6,5		20:00	1,9		21:30	6,8					10 ^a -8-8-8-12 ^a
11/10	07:30	5,3		10:00	6,3		11:30	6,0										21:30	5,7						
12/10	06:30	11,4					12:30	12,3										22:00	6,5						

My data:
- Glucose table (selected)
- Video
- Overview
Information:
- diabetes

Entries: time, glucose value and insuline dose

Daily comments

Screenshot 2

eDagbogen.dk eDiary.dk

Record and play videos

Video playback in browser

Previously recorded video

Change description, view or delete recorded videos

My data:
- Glucose table
- Video (selected)
- Overview
Information:
- diabetes

Figure 2. Screenshots from the eDiary.

Second, the eDiary allows the user to record and play videos directly from the web site with the use of Adobe Flash. By means of the eDiary and a webcam,

² During the pregnancy, six mmol/l is the recommended blood glucose average (Jensen et al. 2004).

video recordings of the women's consultations are easily produced by the women in the outpatient clinic and played at home.

Finally, a collection of links about specific topics relevant for pregnant women with diabetes is accessible from the web portal. The links were selected in collaboration with the healthcare providers who were also encouraged to add links throughout the pilot test. The links were made available through the commercial service, Delicious.

Figure 2 shows screenshots from the prototype. On Screenshot 1, the services (blood glucose, video and links) are listed in the left column. The right column shows an excerpt of the table listing the different glucose measurements along with a timestamp and insulin doses. The last column allows for comments, which this woman used to summarise her daily insulin doses. The recording of blood glucose and insulin intake is focused on predetermined key points around the meals in line with medical advice and the traditional diabetes book. Screenshot 2 shows a list of recorded videos along with a short editable title – pressing the title of the existing video will play the specific video. Furthermore, by activating the top link, new videos can be recorded in the browser from this screen within the eDiary, making recordings of consultations or home video easy for the women.

The main objectives of the eDiary were to support the women in their everyday life and to allow for tele-consultations between the woman and the healthcare providers. Accordingly, choosing a web-based approach allowed the pregnant women with diabetes and their healthcare providers to access the eDiary from any computer without installing extra programs (except Adobe Flash Player for watching video). The system runs and has been tested on all major browsers and platforms. The user interface is developed using Google Web Toolkit Framework. Furthermore, a mobile version is developed to ensure mobility. The Nokia Widget Framework is used to present a compact version of the eDiary on a mobile phone. On the Nokia mobile phone a service makes it possible to enter blood glucose values and also access daily overviews.

While the prototype is aimed at pregnant women with diabetes, the overall architecture is designed to allow the plug in of different services to the system. A modular service-oriented approach was selected to ensure that the diary could easily be reconfigured to support other types of health problems by adding new or removing existing services. In addition, the architecture is highly distributed enabling various vendors to develop different services for the eDiary.

The pilot test

To test the eDiary prototype three pregnant women with diabetes, Martha, Vicky and Emma, were recruited voluntarily in the waiting area at the outpatient clinic. They had type 1 diabetes, Emma only for about one year, the others for more than 10 years. Emma was in her late twenties, Vicky and Martha in their early thirties. The test lasted one month and was designed to support the pregnant women in

already existing routines and activities related to the management of their diabetes. Rather than utilising the traditional diabetes book the women used the eDiary to key in their blood glucose value, the amount of insulin, and comments. During the test the women on average had seven daily entries using either the web page or the mobile phone. They did this either during the day or at the end of the day depending on their other activities and engagements.

The women attended the outpatient clinic every two weeks as part of the ordinary treatment and hence twice during the test (see Figure 3). These consultations were video recorded in the eDiary by the women. In the pilot test, a separate computer was in most cases used at the hospital to access the eDiary to avoid using the IT-infrastructure of the hospital. Emma and Vicky each saw their recordings with their husband once, Martha saw her recordings several times on her own, once with her husband, but faced technological difficulties when showing them to her mother. Vicky showed her recordings to her mother.

The eDiary was used in all but one consultation (due to initial hesitation of the diabetes doctor) with the diabetes doctors to discuss the blood glucose values. The other healthcare providers also had the possibility to look at the blood glucose values, but this only happened a few times which corresponds to the frequency in which they would have used the regular diabetes book. In addition to the ordinary treatment, each woman had one tele-consultation with a diabetes doctor using a Skype video application. Each woman got a scheduled time the day before a check-up at the hospital, but was free to decide from where to have the tele-consultation. A support hotline and support e-mail were available during the test. The women received an eDiary manual and a web camera.

		<i>Obstetrician</i>	<i>Diabetes doctor</i>	<i>Mid-wife</i>	<i>Scanning</i>	<i>Dietician</i>	<i>Tele-consultation</i>
<i>Martha</i>	13-Nov-08	X	X				
	26-Nov-08						X
	27-Nov-08	X	X		X		
<i>Vicky</i>	13-Nov-08		X	X			
	26-Nov-08						X
	27-Nov-08	X	X		X	X	
<i>Emma</i>	13-Nov-08		X				
	26-Nov-08						X
	27-Nov-08		X	X			

Figure 3. Overview of the ordinary consultations (13th and 27th Nov.) and tele-consultations (26th Nov.) during the pilot test.

To collect data we used three different methods during and after the pilot test. First of all, the interaction with the system was logged and data concerning which device was used and when were recorded. Secondly, we wrote extensive field notes and had follow up conversations after each visit to the outpatient clinic. Thirdly, each of the pregnant women, and in two cases also the husband, was

interviewed after the test period. Also two diabetes doctors, one dietician, and one obstetrician who all had been involved in the test were interviewed. All interviews were recorded, transcribed and later analysed with the other data. In the following section, we will present indicative findings from the pilot test.

Findings from the eDiary pilot study

The findings we present here are from a limited pilot study with three women lasting one month. The findings are hence indicative rather than definitive. The starting point of the analysis is the three main services of the eDiary: 'blood glucose', 'video recordings' and 'tele-consultation'. Vicky, Emma and Martha, the three pregnant women with diabetes, were the main users as well as the healthcare providers who treated them. The three women described themselves as experts in diabetes management and successful in managing their diabetes both prior to and during the pilot test. Whether the findings would also be representative of pregnant women who experience difficulties in the management of their blood glucose remains to be tested. The women did not consider themselves expert users of technology. Martha was a clerk and despite her daily use of computers she had only little interest in technology. As a school teacher Vicky also had basic knowledge of computers. She had only little interest in technology, and her husband was the system administrator at home. Emma, who was a PhD student within the field of archaeology, was a more confident user of computers and other technological devices that she used almost on a daily basis.

Using the eDiary to manage blood glucose

Vicky has had diabetes for 11 years, she is married and expects her second child. She is successful in keeping her diabetes tightly regulated and experiences that the eDiary constitutes a useful tool in this. She types in her blood glucose numbers on the mobile phone when she makes the measurement, but feels that the web solution provides a better overview.

Integration in everyday life and work routines

Blood glucose management was a central activity for the pregnant women with diabetes. Figure 4 shows how the women and their relatives used the eDiary at home. Vicky described how she used the eDiary to enter and access her blood



Figure 4. The use of the eDiary by the women in their homes.

glucose figures. Usually she used her diabetes book to write down the figures, but Vicky explained that during the pilot test it quickly became obvious that it was easier to bring her mobile phone than the traditional book and a pen: *'you always bring your phone anyway'*. Likewise, Martha preferred the mobile phone and both agreed that the web portal was excellent for getting an overview. In contrast, Emma preferred to enter the figures via the web solution: *"I'm online everyday so it makes sense to use it"*. She had only used the mobile phone on a few occasions, e.g. when she went away for a weekend, making new entries easy despite being out of daily routine and away from her computer. During the pilot test Emma experienced a change in her need for insulin and used the eDiary in the process of adjusting the dose: *"It's very smart that you can colour the numbers above and below a certain value so you can see if there is a system. (...) During the period where I had to take more insulin it was very pedagogic that I could see exactly where it went wrong"*. Vicky also expressed that the eDiary gave her a feeling of security; *"I don't have to bring my book and if I have forgotten it, then it's just there. And they [the healthcare providers] can find it [blood glucose list] even if I'm not there"*. None of the diabetes doctors had preferences on whether to utilise the traditional diabetes book or the eDiary when treating the pregnant women with diabetes. However, it was crucial that the eDiary could provide them with the standardised overview, as it would be highly time consuming to decipher various systems.

The women's experiences with the eDiary indicate that it integrates well with existing routines, both at the outpatient clinic and in the everyday lives of women. The three women found the eDiary easy available and had each adopted it in a manner compatible with routines of their everyday life and working patterns.




Figure 5. Left picture: Using hospital computer the doctor controls the keyboard and mouse, leaving Vicky as a spectator. Right picture: Martha logs on the dedicated computer to use the eDiary in the consultation. Behind the dedicated computer is the computer of the healthcare provider.

The eDiary at the hospital

The collaboration between the women and their diabetes doctors centered around blood glucose levels, and in the pilot test eDiary served as a collaborative object. Two different approaches for accessing the eDiary were tested. During most consultations a dedicated laptop with a 3G connection was used. Due to problems with the 3G Internet connection the hospital computer was used instead in two

instances. The two different ways of accessing the eDiary web portal produced - through collaboration and negotiation between the women and the diabetes doctor - two different usages (see Figure 5). With the use of the hospital computer the diabetes doctor got the username and password from Vicky, thus gaining control of the situation: deciding when to look at the blood glucose and when to shift to hospital systems, such as the laboratory system. In this situation the women lost control in comparison with the traditional diabetes book where they themselves held the book and could point out important figures.

In contrast, using the dedicated laptop the women themselves logged in and navigated to the appropriate web page with the diabetes doctor as a spectator who should negotiate with the women to access relevant data. Although the situation does not appear to be much different from consultations where the women bring their traditional diabetes book, introducing a new technology opened up for negotiations of the structure of the consultation in the pilot study. A diabetes doctor explained that she often performs various tasks simultaneously and that she preferred to be in control of which tasks to carry out: *“I choose the blood samples and I choose to look in the record or I choose to look at the blood glucose figures”*. However, giving the women a dedicated technology of which the diabetes doctor had no control interrupted the traditional structure where the diabetes doctor sets the agenda, opening for negotiation of the structure of the consultation and the position of both the diabetes doctor and the women. Both Vicky and her diabetes doctor agreed that the eDiary in the pilot test belonged to Vicky. Vicky says, *“it’s my numbers and it is I who can help interpret them – there is a story behind these numbers”*. Her doctor agreed, *“when you access the eDiary then you are on the patient’s turf (...) it’s something we are given permission to look at”*. The concept and design of the eDiary highlighted that the data originates in the home and thus belongs to the women.

Video recording of consultations

Since her first pregnancy Martha’s husband has gotten a new job, making it more difficult for him to attend the consultations of this second pregnancy at the outpatient clinic and thereby share the responsibility. During a single visit to the outpatient clinic, Martha has several consultations, and receives much information, which she often finds difficult to remember.

The pregnant women with diabetes were to a large extent carriers of information both between different healthcare providers at the outpatient clinic, and between daily life and hospital. Being able to remember and to incorporate all information was a critical task to ensure a healthy pregnancy and to feel secure. A day at the outpatient clinic typically consisted of appointments with several healthcare providers making it difficult to take in all information. Martha explained, *“You don’t store all information, only the most important things. But who knows, maybe some of the things you didn’t store could be important too”*. While watching one of the video recordings with her husband and two researchers,

Vicky realised that she was not able to remember everything: *“did she say 3,600 to 3,700 gram? I didn’t tell you [husband] that (...) I am totally surprised”*. The pilot test indicates that through the recordings, the eDiary can support pregnant women with diabetes in encompassing and remembering much information.

For Martha, the recordings improved her husband’s ability to participate in the consultations despite the shift of time and space facilitating their sharing of responsibility. To Vicky and her husband, the recordings improved their exchange of information, her husband explains: *“because when I ask you... it’s always the same to you, so you tell me the same things always. And you can’t remember even half of it. So it’s at good thing to be able to see what happened”*. The healthcare providers also pointed to the potential positive effect of a recording in making the pregnant woman aware of what the healthcare provider really said rather than what they thought he said thereby reducing the level of uncertainty. The women and their husbands agreed that the most interesting recordings were consultations with obstetricians, dieticians, midwives, and scannings as they centered on the baby and provided information new to them.

The pilot test points to the potential of recordings as a supportive tool for both pregnant women with diabetes and their husbands.

Responsibility and system administration

The experiences of the pregnant women with diabetes and their husbands were that both sound and picture should be recorded, that all participants in the consultation should be seen and that gestures should be visible. As the place of a consultation may move from desk to couch and back again, the equipment should ideally be flexible to allow the filming of this automatically.

During the pilot test the healthcare providers reserved the right to refuse to be filmed. On several occasions the healthcare providers discussed the risks of being filmed and perhaps subsequently criticized in public. Most often they came to the conclusion that the advantages were bigger for the pregnant women of having the recordings than the risks they as providers faced being in a Danish context with no or only little tradition for running lawsuits against doctors. However, they came to this conclusion within the context of a pilot study and thus these legal aspects should be further discussed and examined.

Overall the recording of consultations at the outpatient clinic was to the pregnant women’s benefit, raising the question whether the women should also become system administrators, e.g. activating recordings and responding to software updates, rather than the healthcare providers. However, the women are already focused on and engaged in what happens at the consultation and handing over the responsibility to them may be experienced as an extra burden. On the other hand, the women might accept this extra work as it empowers them, This discrepancy between being recorded and getting the benefit is a challenge to the success of using video recordings in this setting (cf. Grudin 1989). In a related

project on video recordings of surgical rehabilitation Sokoler et al present explicit interaction as a way of sharing the responsibility of setting up the consultation and making it explicit when something is being recorded (Sokoler 2007).

The role of tele-consultations

Emma is pregnant with her first child. The frequent visits at the outpatient clinic interrupt her busy workday. She insists on not letting her disease control her and her husband's life. Emma is open about having diabetes and does not consider it a problem to have a tele-consultation with the diabetes doctor from her office that she shares with a male PhD student.

During the pilot test all three women had one tele-consultation with a diabetes doctor. Potentially, tele-consultations can save much time on transportation for the women, and while others have provided larger tests (see (Verhoeven et al 2007) for a literature review), this small scale experiment was set up to explore the outcome and implications of carrying out such consultations in the context of the eDiary.

Emma experienced a delay in the doctor calling her for the tele-consultation, but waiting at her desk she could continue her work. Emma and her diabetes doctor used the web-application of the eDiary to exchange information of blood pressure and blood glucose in the tele-consultation. As they could both see the figures, they were able to discuss them as they would at the outpatient clinic. Since her last visit at the outpatient clinic Emma had experienced a sudden increase in insulin need and was reassured by the diabetes doctor that she had made the right adjustments.

Changing the setting of the consultation

In line with other pregnant women with diabetes who had a stable blood glucose level, Emma thought of the consultations with diabetes doctors as trivial. Rather than getting all the answers from an expert she felt that she and the doctor had discussions where they both had an equal saying. Emma, as with the other pregnant women at the outpatient clinic, often does not see the same diabetes doctor from one consultation to the next. The consultations with the different diabetes doctors do, however, follow the same recognisable structure, where the same topics are discussed and the doctor takes the initiative, which eases the women's interactions with different doctors. The tele-consultation came to follow the same structure making it easy for Emma to interact with a diabetes doctor she had not met before in a way previously not known to her.

However, Martha preferred the consultations at the hospital; *"I prefer to be face to face with the doctor as it's easier to have a conversation"*. To her, the physical atmosphere of the consultation influenced the flow of the conversation. A healthcare provider also expressed her worries about the change of setting, *"I think there might be a risk that it will not be as quiet as needed. You will not put a*

stop to everything at home to have this consultation. The phone may ring, somebody may ring the bell. All kind of disturbances may happen". Compared with the disturbances of the consultations at the outpatient clinic, the healthcare providers were not in control of them in the case of tele-consultations.

From several observations at the outpatient clinic it is clear that the pregnant women seize breaks in the consultation to ask questions. Emma explained that when *"you see that she finds her dictaphone and is finishing up. Then it is about time to ask your question"*. The three women experienced the tele-consultation to be less calm and shorter compared to the consultations at the outpatient clinic, possibly reducing the opportunity of seizing a break. Martha elaborated that it was not only about timing but also about *"remembering a question while leaving the room"*. By having a tele-consultation, the women risk losing the chance to ask a remembered question while leaving the room. These observations are to be considered when changing the setting with a tele-consultation.

As all patients were not as well regulated as Emma, Martha and Vicky the healthcare providers insisted on the possibility to differentiate among their patients in offering this service. In addition, a tele-consultation should be accompanied by a possibility of having a consultation at the outpatient clinic if either the healthcare provider or the patient experienced a need.

Integrating or disturbing

Emma made the tele-consultation from her work. She said that *"if the diabetes should take up as little room in my life as possible, then I need to do all these [diabetes related] things in the situation I am in"*. Vicky also integrated the tele-consultation in her workday. Being a teacher, she prepared herself for the next day's work from her home while having the tele-consultation. Neither Emma nor Vicky experienced that having a tele-consultation in the midst of their everyday life mattered or influenced the outcome. They did not feel that the technology made their diabetes more dominant but appreciated the integration of their treatment in their everyday life. Martha on the other hand did not want to cross the boundary between work and private life; *"I would not like to sit in front of my colleagues (...) I don't proclaim that I'm a diabetic"*. Having the tele-consultation from her home, she was able to keep her private life and work apart.

Tele-consultations might address some of negative aspects, e.g. transportation and time used, of the trend towards centralisation within the healthcare sector. The women in the pilot study inferred that the tele-consultation did not compromise their feeling of security, a feeling they usually got by the many consultations at the outpatient clinic. The pilot test indicates that the eDiary might be a supportive tool to be used in tele-consultations to facilitate exchange of data between healthcare provider and patient.

The eDiary between home and hospital

While the scope of the study is limited, the pilot test revealed indications of minor changes which, in sum and seen in a larger perspective, point to general discussions important for future work attempting to integrate hospital and home. That the introduction of new technology causes changes in practices and collaboration in a working setting is not a new insight within CSCW (Heath & Luff 1996). In our study, however, we focus on changes from introducing technology that connects two very different settings, that of the home and of the outpatient clinic. In the following, we will discuss how the eDiary facilitated the integration and its effects.

The eDiary as an integrating element

The management of a disease requires much work, not only in response to the physiological unfolding of a disease, but also includes the total organisation of the work done, including the impact on those involved with that work and its organization, what Strauss et al has defined as an illness trajectory (Strauss et al 1997, 8). While the term originates in studies of organisation of work in a hospital setting, we believe that the concept also applies to that of the home. Both the pilot study and the initial field study revealed that to have diabetes while being pregnant required the women not only to do an extensive self-care of monitoring their blood glucose but also that it involved a complex organisation of this work, constituting problematic trajectories. For example one woman had a depression in addition to her diabetes and another had work hours that made it difficult to measure the blood glucose at the advised hours. To accomplish the self-care the women used different strategies and at times also involved their husbands (see Ballegaard & Aarhus 2009). Grøn et al (Grøn et al 2008) introduce the concept of homework to broaden the notion of self-care to include the organizational work embedded in illness trajectories that influence the process and outcome of the self-care and hence that medical advice is not always easily followed outside a medical setting.

Much existing diabetes-related technology increases rather than supports or removes the homework, as it requires the user to type in much information (Danholt 2008). The initial field studies drew our attention to the amount of work in self-care and hence the notion of homework and it was a design principle not to add to the amount of homework of the pregnant women with diabetes. As described earlier, the three women in the pilot test did not experience that the eDiary removed their homework nor that it extended their homework. Rather they experienced the eDiary as a support in doing their homework and a tool to ease the integration in everyday life. In developing healthcare IT for the home with the aim of integrating disease management in everyday life, it is hence beneficial to take into account the concept of homework rather than self-care to include the

non-medical factors and the actual work done to manage a disease in the home and in collaboration with healthcare providers.

The role of technology in concealing a disease

An aspect of the pregnant women's wish to integrate the disease management in everyday life was to reduce the space the disease took up. As is often the case with chronic patients (Robinson 1993), the pregnant women with diabetes did not want their disease to control their life and preferred not to be identified solely through their disease, i.e. as a diabetic. As Martha explained earlier, she did not wish to proclaim to have diabetes. To have a chronic disease is to live with your disease the rest of your life. Alonzo (1979) uses the concept 'side-involvement' to shed light on the space a disease takes up in a person's life. As long as you can keep your disease a side-involvement, it does not govern your other activities and is not the lens through which you see the world. Our early field studies revealed that when not pregnant, most of the women with diabetes experienced their diabetes a side-involvement. During pregnancy, it was more difficult for them to keep their diabetes a side-involvement as they were required continuously to do extensive homework. However, they sought to downplay the role of their disease by using different strategies, e.g. concealing artefacts related to their disease or integrating the diabetes related homework in their everyday work.

An objective of the eDiary was hence to support the women in keeping their diabetes a side-involvement. The means to do this were to support homework and to ensure that the technology could be integrated in everyday life without drawing attention to their chronic condition. In creating the eDiary we thus worked with how technology designed for disease management could have functionalities not related to healthcare. While the test focused on disease management, the concept of the eDiary was to merge different spheres of life while still having the opportunity to keep them separate, e.g. to not be reminded of disease when watching private photos. The eDiary was built on technology already integrated in the lives of the pregnant women with diabetes as well of healthy people and hence did not in itself indicate disease. The eDiary provided the pregnant women with a choice to conceal their diabetes status, and its integration helped them keep the diabetes a side-involvement. The pilot test drew attention to the dilemma of integrating while at the same time concealing, and that a healthcare technology should not only take the everyday life of the future users into account but also the perception and nature of the disease.

Bridging home and hospital

While the eDiary integrated disease management in everyday life, the question remains whether it bridged home and hospital in order to create greater coherence in the pregnant women's lives? Field studies established that hospital and home existed as two different spheres, which had an effect on the women's ability to

integrate everyday life with disease management. On the other hand, the segregation gave them instruments to choose different identities in different situations; at the hospital they were patients while at home they were people. The division also affirmed that treatment occurred on the premises of the hospital; it was the pregnant woman who should adapt their work to the consultation hours and the work done in the home was scarcely acknowledged in the hospital sphere. The aim of the eDiary was thus to address the division between home and hospital by making the solution relevant and available for both patient and healthcare provider.

Healthcare technology often belongs to only one domain, either hospital or home. However, the eDiary transcended the domains by placing itself somehow betwixt and between, as both the healthcare provider and the pregnant woman were supposed to use it even if the women were the primary users. In this sense, the eDiary was a boundary object (Star & Griesemer, 1989) inhabiting both home and hospital, although its use and meaning varied between them. The eDiary, as was the case with the diabetes book, bridged the two spheres by bringing information from the home to the hospital and advice from the hospital to the home. In addition, the eDiary bridged home and hospital in making consultations available from home either through recordings or tele-consultations.

A challenge with the design of a technology that can be used in more than one domain is that the user-group is extremely heterogeneous having different needs and routines in which the technology should be integrated. The challenge is to make it plastic enough to match both groups as well as robust enough to be recognizable by both groups, as characterises a boundary object. The eDiary matched the women's needs better than the healthcare providers'. One of the obstetricians said in an interview that he only delivered information to the eDiary, but that he was not involved in the actual use. It could prove to be a weakness of the eDiary as the acceptance and use of a technology, as Grudin (1989) points out, largely relies on the users' ability of seeing benefits in it.

Through the eDiary, the home sphere was strengthened, not at the expense of the hospital domain but as a supplement to it. Neither the pregnant women nor the healthcare system had in this case any interest in abandoning the hospital treatment. Instead, the eDiary augmented the treatment increasing its flexibility of moving between home and hospital.

Structure and hierarchy within the healthcare sector

The healthcare sector today is based on a power relationship that to a large extent is asymmetrical in its structure as it is the healthcare system and providers that set the agenda for the treatment and hence treatment is delivered largely on their conditions. Both patient and healthcare provider recognize their roles and play their role ensuring the status quo of the situation. The asymmetrical relationship is seldom questioned as both parties take it for granted and hence are not conscious

about it or its possibility of being different. However, as argued by Bardram et al (2005) changes may occur in this underlying power structure by the introduction of new healthcare technology. In their study, tele-medical solutions produce new practices, which change not only the communication between healthcare provider and patients, but also the division of work between the two parties where knowledge is collected and interpreted. Similarly it has been argued that to move technology into the homes questions the power relation between clinician and patient and reconfigures the role of being an expert (Ballegaard et al 2008).

The pilot test of the eDiary indicates that the introduction of the eDiary might introduce changes in the practices concerning the treatment that potentially open for a re-negotiation of the underlying power structure within the healthcare system. In our analysis we described how the eDiary offers a new physical space of treatment, new treatment technology, and an empowered patient role: Changing the physical space through the tele-consultation might question the asymmetrical power relation as a consultation from home left the diabetes doctor with little possibility to control the situation as he could in a consultation room and even opened up for disturbances affecting the consultation. Furthermore, being on one's home ground might increase the patient's self-confidence. The possibility of watching recordings of consultations was experienced to be an empowering tool of the patient, as she got the chance to improve her knowledge through seeing the consultation again. Additionally, the women were given the opportunity to question the healthcare providers if she found contradictory information in the recordings. The healthcare providers on the other hand might be more thorough in their utterance as they knew that it could be reheard and discussed at home. The recordings may in extreme cases change the structural power relation drastically as the patient may distribute recordings and use them for lawsuits as previously discussed. Finally, bringing in new technology at the outpatient clinic, over which the women had control in the shape of the necessary passwords and data ownership initiated a potential re-negotiation of the situation.

The structure and hierarchy did not change substantially during the limited pilot study. Nevertheless, the eDiary prompted new routines at the outpatient clinic as the pregnant women with diabetes became users of technology in the consultation rooms and as consultations were also made outside the hospital. The eDiary could provide the healthcare providers with the possibility to organise the work around pregnant women with diabetes in a new manner that to a larger extent could accommodate the wishes and needs of the women. While we acknowledge that the present study is too limited to give any firm conclusions, we find the possible re-negotiation of the underlying power structure to be of vital importance to future work in the design of healthcare solutions that connect home and hospital. It is thus something to be studied further as the implications may hold great potential for rethinking the structure of healthcare services in the future.

Conclusion

Through the design, development and pilot test of the eDiary we have explored effects of introducing technology that supports pregnant women with diabetes in their everyday life. In particular, we have explored the emergence of new practices related to the use of the eDiary and have discussed how new healthcare technology can serve to support patients in the management of their disease in everyday life, and how the introduction of new technology has the potential to open a re-negotiation of the underlying asymmetrical power structure within the healthcare sector.

While the pilot study and the complexity of the eDiary was limited, the study revealed how moving treatments from one setting to another opens a more complex discussion about homework, power relations, different interest in the design of healthcare technology and the challenge of designing and fitting the technology to the everyday life of both healthcare providers and patients. In the case of the eDiary, questions emerged regarding the future of a system, which tend to support and favour the patient and not the healthcare provider, most obvious in relation to the recording of consultations which not only expose the performance of the provider but also is to be used exclusively by the patient. These questions remain open but are highly relevant for future work.

Acknowledgements

We will like to thank all the involved women with diabetes and the healthcare providers for participating in workshops and for testing the eDiary prototype. We will also like to thank Logica and Polycom for participating in the project. Special thanks to Jane Clemensen, Morten Kyng, Lisa Wells, Carsten Munk and Tobias Christensen for participating in the project and ISIS Katrinebjerg for funding.

References

- AccuCheck: (Accessed on March 1st 2009) http://www.accu-check.com/us/rewrite/content/en_US/2.4:40/article/ACCM_general_article_2549.htm
- Alonzo, A. (1979): 'Everyday Illness Behavior: A Situational Approach to Health Status Deviations'. *Social Science and Medicine*, 13A, 4, June 1979, p. 397-404
- Ballegaard, S., Aa., Aarhus, R. (forthcoming 2009): 'Teknologiers mellemkomst i ambulant behandling og egenomsorg' [The intervention of technologies in ambulant treatment and self-care: the case of pregnant women with diabetes], in *Tidsskrift for Forskning i Sygdom og Samfund* [Journal for research in sickness and society].
- Ballegaard, S. Aa., Hansen, T. R., and Kyng, M. (2008): Healthcare in everyday life: designing healthcare services for daily life. In *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems*. CHI '08. ACM, New York, NY, 1807-1816.

- Bardram, J., Bossen, C., and Thomsen, A. (2005): Designing for transformations in collaboration: a study of the deployment of homecare technology. *Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, pages 294-303, New York, NY
- Clausen, T.D., Mathiesen, E., Ekbom, P., Hellmuth, E., Mandrup-Poulsen, T., Damm, P. (2005): 'Poor pregnancy outcome in women with type 2 diabetes', *Diabetes Care* 28(2), Feb. 2005, pp. 323-328.
- Danholt, P. (2008): *Interacting Bodies: Posthuman Enactments of the Problem of Diabetes*, Roskilde Universitetscenter, Roskilde
- Diabetes Logbook: (Accessed on March 1st 2009) <http://www.nesfield.co.uk/diabeteslogbook/>
- Diabetes pilot: (Accessed on March 1st 2009) <http://www.diabetespilot.com>
- Dinesen, B. (2007): Implementation of telehomecare technology: impact on chronically ill patients, healthcare professionals and the healthcare system, *International Journal of Integrated Care*. 2007; vol. 7.
- Dourish, P. (2004): What we talk about when we talk about context. *Personal Ubiquitous Comput.* 8, 1, Feb. 2004, pp. 19-30.
- Grudin, J. (1989): 'Why groupware applications fail: problems in design and evaluation', *Office: Technology and People* 4 (3), pp. 245-264.
- Grøn L., Meinert, L., Mattingly, C. (2008): 'Kronisk hjemmearbejde. Sociale håb, dilemmaer og konflikter i et hjemmearbejdsnarrativ i Uganda, Danmark og USA' [Chronic homework. Social hope, dilemmas and conflicts in a homework narrative in Uganda, Denmark and USA], in *Tidsskrift for Forskning i Sygdom og Samfund* [Journal for research in sickness and society], Nr. 9, pp. 71-95.
- Heath, C, Luff, P. (1996): Documents and Professional Practice: "bad" organisational reasons for "good" clinical records. *Proceedings of the Conference on Computer Supported Cooperative Work*, 1996, p. 354-363
- Jensen, D.M., Damm, P., Moelsted-Pedersen, L., Ovesen, P., Westergaard, J.G., Moeller, M., (2004): 'Outcomes in type 1 diabetic pregnancies: a nationwide, population-based study'. *Diabetes Care*, 27(12), Dec 2004, pp. 2819-2823.
- Lauenborg, J., Mathiesen, E., Ovesen, P., Westergaard, J.G., Ekbom, P., Molsted-Pedersen, L., (2003): 'Audit on stillbirths in women with pregestational type 1 diabetes'. *Diabetes Care*, 26(5), May 2003, pp.1385-1389.
- Mamykina, L., Mynatt, E. D., and Kaufman, D. R. (2006): 'Investigating health management practices of individuals with diabetes', in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI'06, ACM, New York, NY, pp. 927-936
- Mamykina, L., Mynatt, E., Davidson, P., and Greenblatt, D. (2008): 'MAHI: investigation of social scaffolding for reflective thinking in diabetes management', in *Proceeding of the Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems* CHI '08. ACM, New York, NY, pp. 477-486.
- Mendosa: (Accessed on March 1st 2009) <http://www.mendosa.com/software.htm>
- Ministry of the Interior and Health (2003) *Handlingsplan for diabetes* [Strategic plan on diabetes]. (Accessed Oct. 31st 2008) www.sum.dk/publikationer/diabetesbeh/handlingsplan.pdf
- MyNetDiary: (Accessed on March 1st 2009) <http://www.mynetdiary.com/>
- OneTouch: (Accessed on March 1st 2009) <http://www.onetouchdiabetes.com/ultrasmart/index.html>
- Robinson, C. A. (1993) Managing Life with a Chronic Condition: The story of Normalization. *Qualitative Health Research*, vol. 3, no. 1, Feb. 1993, 6-28
- SiDiary: (Accessed on March 1st 2009) <http://www.sidiary.org/>

- Sokoler, T., Löwgren, J., Eriksen, M. A., Linde, P., and Olofsson, S. (2007): Explicit interaction for surgical rehabilitation. In *Proceedings of the 1st international Conference on Tangible and Embedded interaction*. TEI '07. ACM, New York, NY, 117-124.
- Smith, B. K., Frost, J., Albayrak, M., and Sudhakar, R. (2007): 'Integrating glucometers and digital photography as experience capture tools to enhance patient understanding and communication of diabetes self-management practices', in *Personal Ubiquitous Comput.* 11, 4 (Apr. 2007), pp. 273-286.
- Star, S., L. & Griesemer, J. R. (1989): 'Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39', in *Social Studies of Science*, 19, pp. 387-420.
- Strauss, A., Fagerhaugh, S., Suczek, B., Wiener, C. (1997): *Social organization of medical work*. Transaction Publishers: New Brunswick, New Jersey, USA
- Verhoeven, F., Gemert van, L., Dijkstra, K., Nijland, N., Seydel, E., and Steehouder, M. (2007): The Contribution of Teleconsultation and Videoconferencing to Diabetes Care: A Systematic Literature Review. *Journal of Medical Internet Research*, 9 (5). e37.

PRODOC: an Electronic Patient Record to Foster Process-Oriented Practices

Federico Cabitza, Carla Simone and Giovanni Zorzato

Università degli Studi di Milano-Bicocca, 20126 Viale Sarca 336, Milano, Italy
cabitza,simone,zorzato@disco.unimib.it

Abstract. The paper presents PRODOC, an Electronic Document System that allows users to navigate documental artifacts according to predefined process maps. In fact in PRODOC, process models are to be considered as maps that users willingly take as guide for their decisions and actions, rather than scripts prescribed from above. The main tenet of this research is that, by integrating documents and processes, documental practices and related work practices could better align to intended models of action. The underlying concept is the result of a long empirical research in the healthcare domain, where we have deployed PRODOC as an innovative and process-oriented Electronic Patient Record. The user participation in the phase of document definition and clinical processes modeling is central in our approach and it is illustrated in three scenarios of the software informal validation that we present in this paper.

1 Health records: a challenging domain

The healthcare domain is the marketing target of many vendors that propose systems of various kinds to support different phases or activities of the healthcare process. Some of these systems receive a good acceptance since they mainly deal with administrative aspects, e.g., reimbursement accounts, or with quite streamlined and standardized activities, e.g., laboratory examinations. Other systems encounter a stronger opposition since they involve aspects of the clinical care that deal with subtle nuances of clinical work: among these systems there is still the Electronic Patient Record (EPR). EPR introduction in the hospitals is still rare and in any case problematic (Heeks et al., 1999; Hartswood et al., 2003; Berg and Winthereik, 2003;

Koppel et al., 2005). This was also manifest during our interactions with several practitioners working in hospitals that were at different stages of introduction of EPR systems. We noticed a diffuse sense of frustration in these clinicians towards this process, which in many cases was almost out of their control. In fact, decisions about the digitization of patient records (PRs) were often governed either by external forces, like the bargaining power of influent vendors, or by internal strategies, which were aimed at a uniform adoption of the same system across different departments to improve cost control and resource management.

Obviously, for us as external researchers, these processes were out of our influence. We had nice conversations with the practitioners about a matter that they must have had at heart, and we met their open availability to validate some of our ideas on small prototypes and mock-ups in sort of “loud thinking” users sessions. This gave us the motivation to capitalize the wealth of experience gained with practitioners and try to address the point whether a really *innovative* EPR could be possible. We agreed with practitioners that innovativeness was all about the challenge of fulfilling their primary needs (i.e., care and its account) without requiring them to distort their usual practices and saddle themselves with the low-level integration with the hospital information system. We decided to face this challenge since we believed that our outcome would at least provide practitioners with a tool that they could exploit in their interaction with the hospital management and the ICT vendors to spur them towards real innovativeness. This paper is a first step in this direction: after discussing critical aspects of existing EPRs, it gives the basic tenets and describes the current version of our outcome, a prototype called PRODOC, together with an account of its ongoing validation.

2 Current EPRs: a critical view

Our empirical research about the introduction of EPRs in four hospitals in Northern Italy highlighted three main critical areas that deeply influenced the design of PRODOC:

First: Standard EPRs provide their users with a sequence of electronic forms that reflect how information has been modeled to deploy the underlying database: i.e., in terms of domain entities and corresponding relations. These aggregate views propose layouts of “assembled squares” where practitioners can read or write clinical data according to predefined schema. Moreover, these forms are linked together according to a business logic that, in general, has nothing to do with the paper forms that are in use before the EPR’s introduction and with the practices that doctors have built around them. Several studies (e.g., (Harper et al., 1997; Fitzpatrick, 2000)) pointed to the advantages of the paper-based forms over the electronic ones but the focus was more on the affordance of the paper medium rather than to the internal structure of these forms. Indeed, we believe that how doctors organize information in their records (a task that is facilitated in paper-based forms) is crucial: in fact, we observed that the structure of a document/record/chart is usually the outcome of a long-lasting process where the results of consolidated work practices and conven-

tions have been stratified. In particular, this structure is able to let tacit knowledge be evoked by the mutual position of information or by specific graphic cues or textual annotations that are easily juxtaposed beside or around it. All this rich combination of tacit, implicit and explicit knowledge is mainly lost when EPR is digitalized: the consequence is that practitioners are requested to abandon the practices on which they usually base the effectiveness of the care process, and to behave according to something that is out of their experience: i.e., the business workflow logic of the EPR.

Second: The business logic of EPRs are usually invariant with respect to the specific care processes that doctors are able to tailor to specific diseases and to how patients react to their interventions. These processes can be implicit medical knowledge, or made explicit in what is usually called a Clinical Pathway (CP) (Sloan and Guinane, 1999); the whole body of these processes is a sort of procedural knowledge that in any case clinicians retain and exploit to articulate their actions and those of the practitioners involved in the same clinical case. Generally, in current EPRs there is very little or no support towards these disease-related and patient-centered processes: and no wonder there is not. In fact clinical processes are defined, updated and dismissed not only according to medical evidences and guidelines but also according to very local drivers, like available resources, staff, level of education, available equipment, and even hospital topology (Lenz et al., 2007). Therefore, in the more positive case, the burden to “remember and follow” the intended pathway is left on the shoulders of practitioners; in the more negative case, they have the additional burden to cope with a contrasting logic embedded in the EPR.

Third: In specialist literature, doctors themselves propose the CP as an effective tool to decrease undesired practice variability and improve clinical performance (Campbell et al., 1998). We also, in Cabitza et al. (2008), observed the practice of adding the sheets of a reference CP into the record folder of a patient with a specific disease. In those settings, and especially when the CP is the outcome of a bottom-up collaborative effort, the integration of CPs – however represented – and the EPR is advocated to improve appropriateness and to aptly respond the increasing demands for patient safety, better risk management and reduced costs. In fact, studies have shown that IT applications can increase pathway compliance, if they embed pathways in routine work, and more precisely in routine documentation (Lenz et al., 2007). To this aim, a traditional approach is that of conceiving a set of electronic checklists that allow doctors to check the compliance of their practice against the pathway (Blaser et al., 2007). Yet this approach, even if mitigated by the principle of “charting by exception” proposed by Short (Short, 1997), usually results in documental overhead because clinicians end up by reporting more *about* the pathway than *according to* the pathway. As discussed by Berg (1999), clinical reporting follows clinical work closely and clinical work is influenced by how and when clinicians report it since they rely on the record to coordinate with each other. Even when EPRs acknowledge this mutual influence, they embed processes that are interpreted as “yet other” workflow models and scripts by “engines” that govern what tasks the application can allow users to perform, or as the outcomes of domain

knowledge representations by inference engines oriented to planning and decision making (Smart and Roux; Quaglini et al.; Aigner and Miksch, 2006).

3 Innovative EPRs: the basic tenets

To overcome the above limits, we identified basic tenets that we adopted for our proposal and express in what follows as they were formulated in frank terms by practitioners:

“Let me keep my folders!”: from their introduction at the end of the 19th century, paper based PRs have evolved in very well organized bundles of documents, charts and records. What to software engineers could seem confused folders are actually sets of sheets that are grouped together according to their contents, to the time span they refer to and to the phase of the care process they are associated to (Cabitza et al., 2005). Doctors see this whole information not as a static “folder” but rather as a “web of artifacts” (Bardram and Bossen, 2005) whose ad-hoc and often unpredictable organization allows for different levels of aggregation and retrieval; this flexible folder allows doctors to continuously re-arrange its sheets so that, e.g., the peculiar proximity of pieces of information can facilitate peripheral readings that profitably complement data that are on current focus; or it can allow for a comprehensive view of data according to different criteria, like time intervals or basic indicators of patient’s condition (Fitzpatrick, 2004). In the EPR design, this means that the “folder” metaphor by which to gather and present clinical data must be preserved against the omnipresent metaphor of the “dashboard”, borrowed by other information-intensive domains.

“Let me do what I do on paper”: our observational study confirmed that practitioners appreciate the possibility to keep browsing and skimming clinical data as they were used to with paper-based bundles as well as to add extemporary and informal annotations (Hardstone et al., 2004; Bringay et al., 2006) to the basic structure of documents: more specifically, practitioners claimed that such flexible annotations are a natural means to promote awareness and to evoke tacit knowledge pertinent to the annotated information. Practitioners also appreciated the possibility to customize, tailor and design their own forms according to their local needs as they were used to with paper-based forms, which they usually could compose and print locally. These continuous improvements always require a formal validation by the hospital management but they are on ordinary agenda if doctors comply with the core data set that has been defined at organizational level. Conversely, modifying an EPR requires much more than mere negotiation with the management since it always requires modifying the EPR DBMS and often even its internal business logic.

“Integrate data and processes but don’t mix’em up”: usually documents and processes are seen as independent units or, better yet, able to characterize a work domain at different levels. However, clinicians told us that when a representation of an organizational process is concerned with the clinical dimension of hospital work, they should be seen as simple maps, rather than “scripts” (Schmidt, 1997) of an application logic that prescribes and steers clinical behaviors. The term map

suggests that these representations can be used to provide a sort of loose “topology” that can promote awareness of the unfolding of the illness trajectory, of what activity is currently being performed and of what doctors should do next to coordinate with each other in a seamless way. The process can then be seen as an alternative way to organize data in terms of *when* they have to be produced and to *what* aim. Therefore, to improve awareness and coordination, clinicians suggested that there is no need to build a comprehensive model of the care process, whereas it suffices to make the crucial input/output relations between activities and specific documents (or parts of them) clearly explicit (Cabitza et al., 2007). Moreover, in the healthcare domain, process maintainability is a pressing requirement, as we reported in Cabitza et al. (2008) since strict standardization from above built into electronic systems is bound to fail (Berg, 1997); indeed, for clinical pathway to be effective aids in guiding practice, doctors must be able to continuously update and maintain them, according to both local practices, new scientific evidences and agreed guidelines based on consensus within a particular discipline.

4 The basic design choices

In order to build a prototype fulfilling the above basic tenets, we had to make choices and tradeoffs about how to represent documents and processes. To this aim, we made a survey of the main solutions reported in the literature and collected examples from the hospitals we have been in contact with in the last years. In regard to how represent clinical processes, the most used formalisms are extensions and customizations of the flowchart notation. In a minority of cases, the representation is based on matrixes, or time-grids, describing activities and responsibilities. More often, protocol-based care is formalized in terms of hierarchical networks of component tasks, that unfold over time, like Asbru (Shahar et al., 1998), GLIF (Boxwala et al., 2004) and PROforma (Sutton and Fox, 2003) (just to cite the most used). An alternative approach takes a declarative point of view, e.g., the declarative language CIGDec (Mulyar et al., 2008), to increase the flexibility of process execution: in fact, activities can be modeled without causal relations and possible precedences between them can be expressed in terms of additional constraints. In this rich panorama, our choice was to combine the advantages of an explicit representation of the relationships among activities with the flexibility resulting from releasing this representation from the engine that is usually associated to procedural languages. This suggested us to adopt the Business Process Modeling Notation¹ (BPMN), a standard notation close to flowcharts, since we observed it is familiar to most of the healthcare practitioners for historical reasons

In regard to how represent document structures, we adopted an approach that is common to other applications (see, e.g., Morrison and Blackwell (2009)) that

¹ The Business Process Modeling Notation is a graphical representation developed by the Business Process Management Initiative (BPMI) and currently maintained by the Object Management Group.

try to mimic the look-and-feel of paper-based chart quite closely. More than this, our solution aims to provide an interaction mode that makes the definition of document templates as natural as possible for the practitioners. This solution invites practitioners to see documents as just topological arrangements of data fields or “document constituents” that we call *didgets* (more details will be given in the Section 6.1). Our point is that avoiding any explicit representation of the relations between these constituents allows for a great level of flexibility in presence of changes in the forms’ layout as well as of the insertion/deletion of new pieces of information. This choice has had a strong impact on how to realize the clinical data repository behind our EPR: in fact, its structure is dynamically derived from the pieces of information contained in the document templates by adopting a multidimensional approach (Pedersen and Jensen, 2001).

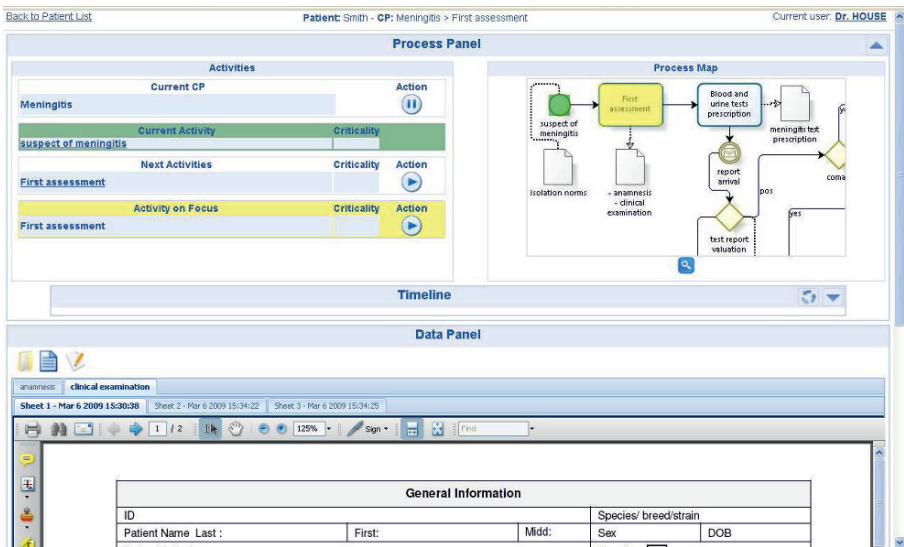


Figure 1. The main page of PRODOC (the Timeline is minimized).

5 PRODOC: description of the user interface

In this section, we outline the main functionalities of PRODOC (PROcess-oriented DOCumentation) a web-based prototype of EPR aimed at supporting clinicians in integrating clinical pathways into their records. First, we describe the main page that PRODOC displays once a doctor has selected an activity from the list of current CPs that are active on a particular admitted patient. This page is divided in two main sections (see Figure 1): the upper pane is called **Process Panel** and provides process overview and navigation functionalities. The lower pane is called **Data Panel** and provides user with access to data through paper-looking documents (currently, they are PDF forms). In its current version, the Data Panel provides also annotation

functionalities through the rich command palette provided by the Acrobat platform. The Process Panel can be reduced (or minimized) so as to give users a full-screen of the Data Panel: this can be particularly useful if PRODOC is used on a tablet or ePaper PC; in fact, this full screen view allows to simulate the traditional interaction with paper-based artifacts in all those settings where this is considered a plus by practitioners. Above the two main panels, PRODOC encompasses also a small textual section, which reports the main patient personal data, the navigation trail and information on the current user.

5.1 The Process Panel

The Process Panel allows users to have a quick glance of the process map, to assess and set the current state of the clinical process, and consult the process history. To this aim, the panel is divided into three sections (see Figure 2). Two of these sections, the **Process Map** and the **Activities** are fixed, while the third, the **Timeline**, is collapsable as the whole Process Panel.

The **Process Map** is a window where a portion of the graphical BPMN-based representation of the current CP is displayed. The map is automatically centered around either the active activity or the activity currently on focus (i.e., the activity that the user has selected to view the associated documentation) but it is also draggable, so that users can examine different regions of the process schema. The process map is an active map: this means that the diagram elements depicted therein are active links that make an activity on focus and its associated documentation be displayed in the Data Panel. PRODOC also provides a *zoom function* (the magnifying glass icon depicted in Figure 2) that allows to enlarge the map and see it in full. In the current prototype, the active activity is highlighted in green, while the activity on focus (when this does not coincide with the former) is colored in yellow. The Process Map works in combination with the Activities section.

The **Activities section** reports textual information about i) what the current activity in the process is (and its criticality); ii) what activity/ies follow/s the current one (and their criticality); iii) what activity is currently on focus: its documents are currently displayed by the Data Panel. In regard to the current activity, users can suspend it, by pressing the `Pause` button. Since the current prototype does not handle parallelism within the same CP, two activities of the same CP can not be active at the same time. Therefore, if an activity is temporarily suspended, also the overall CP is suspended and an event of temporary exit from the CP is generated. When the CP is suspended, it enters a sort of *unspecified activity* where users can get access to (and use) all the patient documentation, The suspension lasts until either a new activity is started or the current activity is resumed. When a user means to terminate an entire process and exit the CP, she can select and make one of the *end event* elements (depicted in the process map) active.

User can select any activity to put it on focus in order to read or write the associated documents. Once an activity has been put on focus (through either the Process Map or the Next Activity link) users can activate it by pressing the `Play` button, to

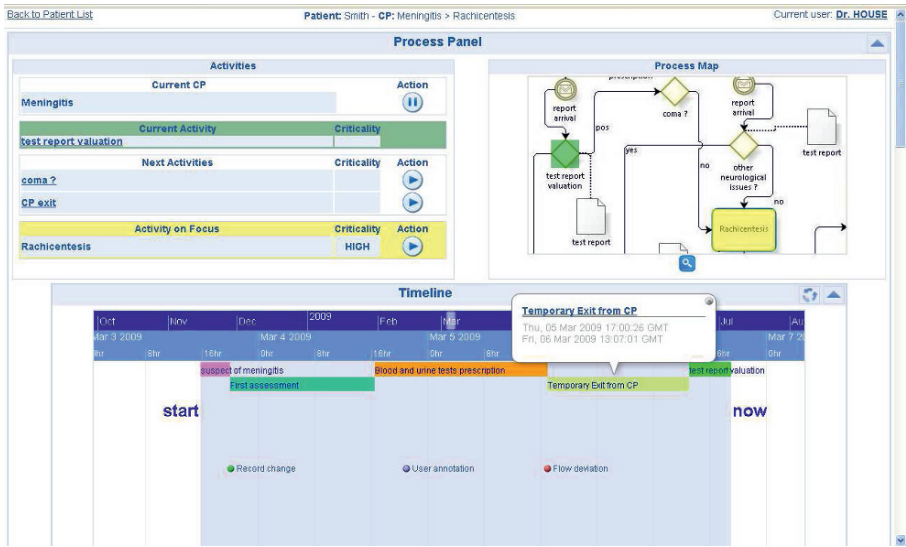


Figure 2. Process Panel with the Timeline open.

tell the system this is the activity they are currently involved in. In so doing, the system automatically terminates the current activity, makes the activity on focus active and updates the next activity links accordingly. If the activity that the user activates does not directly follow the current process activity (i.e., it is not one of the next activities), the system generates a *deviation* event that is recorded to be displayed in the TimeLine.

PRODOC does not pose significant constraints to user action and optionally can ask users to provide a written justification if they start a new activity without following the order indicated in the clinical pathway. This request can be set as either mandatory or optional by the responsible of the CP (i.e., usually who led the task force supposed to design and maintain the CP from the available guidelines) according to local conventions and if this indication is deemed useful to provide doctors of next shifts with the context to understand past decisions; in our specific case, the CP responsible proposed this functionality as a way to foster feedback from doctors on the extent the clinical pathway at hand is useful and reliable in their daily practice.

The **Timeline** is a section at the bottom of the Process Panel that can be expanded and collapsed at need every time doctors need to get a visual representation of when relevant events occurred and during what activity. The timeline displays both the process history, i.e., the sequence of CP activities that doctors have actually performed till the present moment, and any relevant event that has occurred during the patient stay that far. When users open the Timeline, it is centered on the present time if the process is still in progress, or on the end of the healthcare episode, if the patient has been already discharged. When the Timeline is updated, PRODOC displays unobtrusive pop-up messages, so that users can determine whether to refresh the timeline by pushing a specific Refresh icon. In the Timeline, activities are

depicted as a succession of colored bars, while events as dots of different shape and color according to their predefined type (e.g., changes in the record, report arrivals, *deviations* in the CP trajectory). Users can scroll the timeline along the horizontal axis by two graduated scales to explore the process history with different time granularity: the former scale is divided in monthly intervals and allows for a quick shift upon the time axis; the latter one splits the timeline according to days and hours and allows for more accurate movements. When a user clicks on an activity bar, PRODOC shows a pop-up balloon that reports the start time of the activity, the justification given for its activation (if any), its end time (if already performed) and a direct link to the associated documentation and contextual content. This latter functionality means that doctors have got a sort of “time machine” by which to see the record’s content at the selected time: pages accessed through the timeline panel are displayed in the Data Panel as usual but are read-only unless the activity on focus is the current one. In regard to the relevant events displayed in the timeline, the current implementation of PRODOC considers three event types:

- data events (green spots), which are pinpointed into the timeline whenever users insert new data in the Data Panel (i.e., save the content at sheet level);
- annotation events (blue spots), which inform users of when their colleagues have annotated a document.
- deviation events (red spots), which are displayed either when users start an activity that does not follow the current one in the CP map; when users write on a document that is not associated with the current activity; or when users suspend a CP.

If a user selects either a data or an annotation event, PRODOC displays a balloon where exact time of document saving and the author identity are indicated, as well as a direct link to the saved/annotated document.

5.2 The Data Panel

Below the Process Panel, users can see the Data Panel (see Figure 3) in which PRODOC embeds the set of the only documents associated with the activity currently on focus. If, conversely, the user has to consult the whole documentation during a specific activity, she can get access to all the sheets regarding a single patient, by pressing the `Display All` icon (the first from the left in the command bar depicted in Figure 3); PRODOC displays the whole documentation for a specific patient also after the user has temporarily suspended the CP.

In the Data Panel, users can read and write the forms that are progressively compiled during the patient stay. Usually each activity has some document templates and sheets associated; users can swap from sheet to sheet pertaining to different templates by means of tabs, so as to mimic how they are used to in the case of paper-based folders. When users complete a sheet, they can have PRODOC generate a new sheet for each template by selecting the `New Sheet` icon (the second from the left in Figure 3). By affinity with the typical constraints of the healthcare domain where each inscription is a legal report, inserted data cannot be changed nor

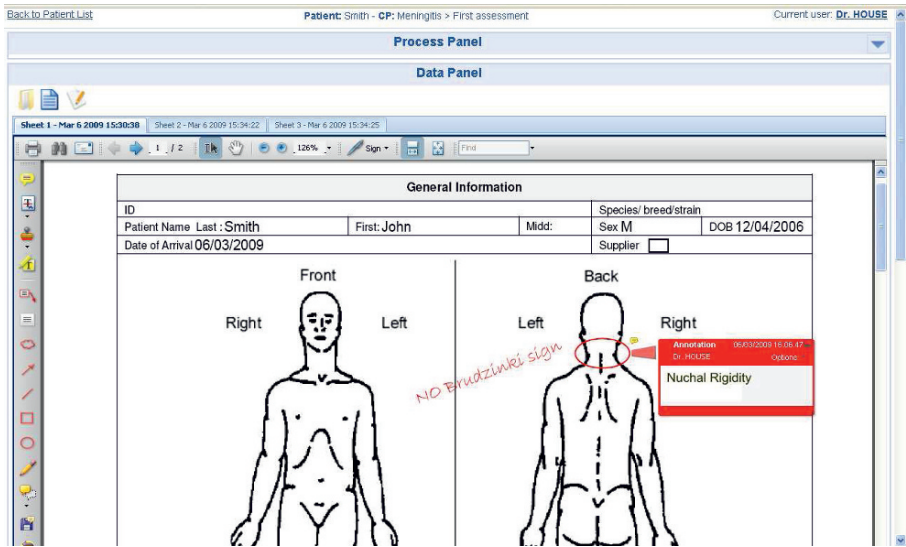


Figure 3. The Data Panel displaying an annotated form.

erased. Inscriptions can be stricken-through and new data can be juxtaposed beside the obsolete (or just wrong) ones as they would on paper².

Sheets are opened in read-only mode. The user that has accessed a sheet first can also turn on the edit mode (third icon from the left) and get a lock on the document for a particular patient. In this way, her colleagues can still consult the sheet in read-only mode but cannot edit it till she has released the lock by either saving the last changes or discarding them. The Data Panel provides user also with versatile annotation capabilities. Since, in the current version of the PRODOC prototype, templates are PDF forms (primarily for their strong resemblance with paper-based forms and the familiarity that generic users have with this format), each sheet is annotatable with the rich palette of drawing markups provided by the Adobe Acrobat platform: users can then add textual comments, either by keyboard or by the free-hand pencil tool; they can draw oval, circles and rectangles around portions of text or pictures and mark arrows, in a very similar way as they use to do on the paper-based record with pencil inscriptions, sketches or with post-it notes, in order to communicate with their colleagues, drop a note to recall in the next shifts (Bringay et al., 2006) and unobtrusively promote collaboration awareness and coordination. Each annotation is registered as a separate event, characterized in terms of time and author, and it is consequently reported in the Timeline; even more noteworthy, as soon as a user adds a note into a sheet, this event is immediately notified to all the users that have the same sheet open, so as to facilitate synchronous communication and make actors aware of any change in the documents that they are reading.

The main point to retain here is that PRODOC can be used in three different

² This constraint can be easily relaxed if PRODOC is to be used in different documental domains.

operating modes: either as a sheer electronic record where any artifact is accessible with no order constraint, if users minimize and disregard the process panel; as an interactive process map, if users have to use another third-party or legacy EPR and still want to keep trace of the current activity in the context of the adopted CP; or as an integrated tool that enables process-oriented document navigation and event-based information retrieval. In the next section, we illustrate the main functionalities of PRODOC in the context of three scenarios, by which we undertook the informal validation of the current prototype.

6 Participatory discussion through use scenarios

From the beginning, we realized that an effective evaluation of PRODOC could not be conducted in a laboratory setting. In fact, PRODOC has been designed to support the ongoing recording of (clinical) acts and the articulation of activities unfolding around this general task. Therefore, a true validation of PRODOC would require a long-term deployment and an analysis of the impact of the system on both clinical and documental practices. On the other hand, to gain an initial feedback on the effectiveness of the main functionalities of PRODOC, we undertook an informal validation according to a qualitative approach that encompasses the involvement of a small user panel. To this aim, we have so far conducted informal evaluations with clinicians in the context of three simplified use cases: one case regarding the definition of the patient record templates (to be displayed in the Data Panel); another case regarding the definition of clinical pathways (to display in the Process Panel); and one case regarding their combined use in a realistic clinical scenario.

These use scenarios were reviewed and personally experienced by selected practitioners during individual user sessions, lasting approximately forty-five minutes: a specially instrumented version of the PRODOC prototype was deployed in the hospital library so as to monitor the browsing activity, command selection and software events triggered by users during application execution. For our user panel, we invited the head doctor, a senior doctor that the former invited for his past experience with EPRs, the head nurse and one of her most experienced colleagues to use the system following the three scenarios as a sort of loose plot. The evaluation methodology we used was based primarily on usage logs and user feedback gathered in approximately one-hour long semi-structured interviews taken immediately after the user sessions. These interviews were used to support our understanding of the usage logs, to acquire feedback about how the tool was perceived and keep track of relevant suggestions. This first round of evaluation sessions, although informal, allowed us to gain insights into how PRODOC would be understood and used by clinicians to get access to their daily documentation and comply with the specifications of clinical pathways. In what follows, we will run through the scenarios we proposed to the user panel and will interpose the main remarks of the involved users while describing in some further detail the main functionalities of PRODOC's current implementation.

6.1 Document schema definition

Before PRODOC can be used, its users have to create both the process schema and the document templates it will use. This phase of preliminary definition is an important part of the innovative approach of PRODOC to process-oriented documentation. In fact, on the one hand, we wanted users be as much independent as possible in creating their own processes (i.e., clinical pathways) and correlating them to their own records. On the other hand, PRODOC also proposes a strong document-centered approach to data, i.e., we wanted users come back to thinking of data as elements of specific and “tangible” documents, rather than aggregated elements taken from underlying databases and gathered together in sort of virtual views.

To this aim, we invited our user panel to use an editor by which to re-build their own document templates and make them as similar as possible to the paper-based charts they were currently using. This was something that two practitioners we involved were already used to: in fact, they were members of the large group of hospital representatives that was supposed to compose (and maintain) the templates of the hospital patient record, have these validated by the hospital management, convert the validated templates into PDF files and then share them to their ward colleagues, so that these could print the blank charts at need to progressively feed their paper-based folders. The template editor we provided to the user panel was an augmented version of a very popular word processor, already used at the hospital: we developed a Visual Basic application that, while the word processor is open on a new document, displays a sort of small floating panel from which users can select the proper *didgets* to insert in the document. We called *didget* a *documental widget*, that is, a coherent set of data fields that is convenient to gather together at a certain level of description. Following the scenario, the user panel was invited to drag a *patient didget* from the editor palette and drop it on top of a chart template in order to add the usual patient data (e.g., name, surname, date of birth) in the chart header; users could also choose a *prescription didget* (which encompasses fields for the drug name, dose, the scheduled administration time, etc.) to create a new row for the table of the drug prescription form. We explained that *didgets* can be any element of a typical form: either simple input fields, or set of fields, multi-line text areas, check boxes, drop-down lists or combo boxes, according to the minimum data set that key users, domain experts and system analysts had collaboratively identified from the domain and document analysis. The user panel saw that the same *didgets* could appear in more templates and, obviously, in more instances of a single template, i.e., in more sheets. In this latter case, a *didget* could be *local*, if the data associated with it are not to be replicated in other sheets; or *global*, if the data must be replicated in each occurrence of the same *didget* all over the record.

Seeing document templates as containers of *didgets*, and the task of document editing as that of positioning *didgets* in a blank page was not a practice that fitted the habitual practices of our panel easily and immediately. Yet, at the end, the concept of *didget* was received quite warmly: users understood its scalability from a single text input field, e.g., a body temperature field, to a complex record, e.g., the matrix of a fluid balance charts, that is a “field macro”, as it was suggestively

called by the head doctor, which could be reused in every record and chart where those data need to be reported and consulted. On the other hand, users realized that the document structure, i.e., how data are displayed in a record, and data fields could be *decoupled*. The senior doctor told us that he worked for a couple of years in a hospital ward where doctors had been using an EPR for years: he told us that after only few days from the first deployment he and some colleagues of his noticed that the body weight field was in the wrong page of the software application, i.e., associated to a preliminary phase of patient admittance where clinical data were not collected; and that the same field was not replicated in another page where it could have been useful to calculate drug dosages. He told us that they asked the software vendor to change the user interface accordingly and that when he moved to the NICU, a couple of years later, they were still waiting for this patch. He told us this anecdote since he realized that in PRODOC he could have used a regular word processor and changed the user interface just by moving (or importing) a didget in a regular document. If this didget had not been anticipated by analysts, it could just have been added to the underlying DBMS as just another dimension related to the patient. In more general terms, the capability of changing the user interface of PRODOC (i.e., what is embedded in the Data Panel) at need and according to very local, specific and ad-hoc needs by just creating new document templates was seen as a clear plus of our approach.

6.2 Process schema definition

Once users had created the templates of their record, we told them that it was time to create their first pathway and correlate its activities with (not necessarily exclusive) sets of templates. In order to create a computational specification of a clinical pathway, we provided users with a graphical process editor that we chose because it is freeware and has got a very user-friendly graphical interface: BizAgi Process Modeler by BizAgi. This editor allows to create process diagrams in the BPMN standard format by simply dragging and dropping iconic elements: boxes for activities, diamonds for decisional branching points, circle for events and oriented arcs for flow relations. We invited the users to characterize the pathway elements by specifying their criticality (i.e., either importance or prescriptiveness) and that with this editor users could specify in terms of an extended property of either activities or flows. To associate activities with the set of templates that could be considered as either input or output of the related tasks, we instructed users to employ the BPMN constructs “association” and “artifact”. Due to the still semi-automatic integration between BizAgi and PRODOC, users had to write the name (URI) of the templates related to each task as a property of the artifact construct. For the next versions of the tool, they advocated a visual mechanism of template importing similar to the didget floating panel. While our panel showed it was proficient at modeling a pathway in terms of activities and branching points, it required some time in getting a clear comprehension of when and how to use the BPMN construct “event”. This is an additional feature with respect to traditional flow-charts, but users acknowledged

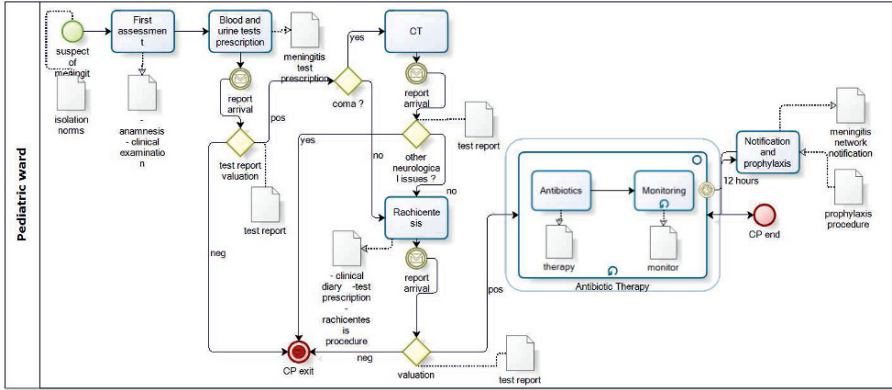


Figure 4. The clinical pathway doctors designed for meningitis cases.

its aptness to report either time- or data-related aspects of CP (e.g., “when the report becomes available”, “after twelve hours”).

In the scenario, after that users have diagrammed a pathway and indicated the relationship between tasks and documents, the editor then automatically exports the process specifications to PRODOC, in terms of both process map (as a raster image) and computational representation (as an XPDL³ file). These resources are then used by PRODOC to enable the document navigation according to the process model.

6.3 Using PRODOC in a clinical scenario

After the two scenarios of process and template definition, we asked clinicians to consider PRODOC as their official patient record and fancy themselves involved in the management of a patient with a suspicion of meningitis. The corresponding CP, drawn by means of the editor mentioned in Section 6.2, is shown in Figure 4.

The user panel had to imagine a situation in which an infant, JD, was admitted in their pediatric ward with a meningitis suspicion formulated by the Emergency Department (ED) on the basis of clinical signs. The responsible physician of the ED had opened PRODOC and associated a *Hospital Stay* meta-process to JD so as to fill in the first pages of the corresponding patient record. A meta-process is just like a regular process but it can also contain sub-processes: i.e., activities that can be further characterized in terms of other processes. In this case, the *Hospital Stay* is a hospital process that encompasses very general activities in strict sequential order: the Patient Admission, the Patient Treatment and the Patient Discharge; the Patient Treatment task, in turn, can be associated with several other sub-processes, i.e., instances of different Clinical Pathways (CPs) by which doctors can manage the health problem identified in the Admission phase till discharge.

³ XML Process Definition Language, standard developed by the Workflow Management Coalition (WfMC) to interchange process definitions between different management tools.

In the Hospital Stay process, the first activity is an Admission, and this is related to the document templates where the hospital Triage records the admission, identifies the patient and reports a first set of possible diagnoses. According to the Triage evaluation, JD is transferred to the pediatric ward. The admitting pediatrician opens PRODOC, selects the JD record, puts the Admission activity on focus and then consults the associated Triage documentation; after that, she starts the Treatment activity. When a subprocess task is selected (i.e., in this case, Treatment), PRODOC prompts for the selection of a subprocess that specifies the general task: in this case, the pediatrician decides to associate JD with the specific CP that the hospital published to cope with admitted cases of suspected meningitis.

In regard to this first part of the scenario, the user panel appreciated that a patient could be associated with several processes (and hence CPs) in parallel: although they acknowledged that they associate more CP to the same patient quite seldom, this would regard a number of complex cases when a patient is transferred from a facility to another for either complications or further investigations while she is still under the partial responsibility of the former facility: an equivocal situation that is usually difficult to manage when it is not clear who can do what. Practitioners also appreciated that a general pathway could include more specific processes: indeed, they claimed that a number of their pathways were quite “abstract” and with activities that often should need to be further characterized in terms of standard procedures, detailed diet regimens or more refined treatment protocols. Yet, at this point, the first of the main shortcomings of the current PRODOC emerged. In fact, the user panel was agreed that even if a patient is managed according to multiple pathways, a number of sheets from her patient record must be *shared* across these pathways, since they could be read and written in activities of any CP: the possibility to use the `Display All` icon (mentioned in Section 5.2) to see all the sheets related (through their father templates) to the running metaprocess (and to all its child processes) was considered a too complicated “trick”.

Next in the scenario, the pediatrician consults the graphical map of the meningitis CP in the main page and that reminds her of what to do first when coping with such a case. The first activities mentioned in the CP are the *First Assessment*, which regards reporting the anamnesis and the findings of the clinical examination, and the *Test Prescription*.

Three users proceeded in this order and activated First Assessment by clicking the `Play` button. The senior doctor said that often, according to subjective impressions, diagnostic tests should come first in order to gather sound evidences for the correctness of the Triage diagnosis. In this case, users observed that the pediatrician could put Test Prescription on focus and then activate it. The point highlighted by the user panel was that PRODOC does not impose a strict order or restriction on when to use a particular document, such as it could happen in a workflow management system; rather, it only suggests users what set of documents they could need while they are performing a specific task, it gives them the possibility to create new sheets, or to open those already existent to accomplish a specific process activity. When the head nurse opened the Test Prescription form, she noticed that the pa-

tient’s details in the header were already filled in: in fact, those data were coming from the global didget “Patient’s Details” of the Admission form that was compiled in the first activity. Every new sheet from a template that has got that didget will have those data replicated with no additional effort from the practitioners.

When the pediatrician selects the activity Blood and Urine Tests Prescription without passing through the First Assessment, the timeline records the event as a deviation from the CP. The Test Prescription activity is associated with a Meningitis Test Prescription document template. The pediatrician marks the tests that she needs for the meningitis diagnosis (e.g., blood count, PCR and glycemia) and then saves the sheet. When the sheet is saved, PRODOC stores the data in the underlying database and records corresponding data events in the process history (timeline). Afterwards a nurse consults the Meningitis Test Prescription sheet, takes the blood samples and marks the checkbox “specimen taken” in the Prescription sheet.

After we had illustrated the part of the scenario involving the nurse, the head physician and head nurse noted that it would be important to explicitly distinguish what different roles are involved in a particular task. This was the second main shortcoming of PRODOC reported in the user sessions. In fact, role management is still preliminary in PRODOC, which currently only distinguishes between the “responsible” of the CP and simple performers. While the latter ones can change any document related to the CP, only the owner can create new process instances, terminate and activate activities or suspend a CP, i.e., change the CP status. As a matter of fact, roles can be easily represented in the process map by exploiting the BPMN elements “pool” and “lane”, used to represent different participants in a process and to organize activities within pools, respectively. Users said that this knowledge about “who can do what” had to be reflected in PRODOC, in terms of transparent management of different read/write access permissions to documents according to the role associated to the user currently logged in⁴.

Coming back to the scenario, while doctors are waiting for the lab to analyze the specimens and issue the report, a pediatrician can decide to prescribe an anti-inflammatory drug. This task is *not* represented in the CP schema because it represents an exceptional decision that really depends on the particular conditions of the patient. Obviously, a CP schema cannot include all the possible exceptions that may occur during a treatment (even if it should consider the most important ones and those that usually lead to aborting it). As a consequence, PRODOC allows for writing additional documentation that is not related to any activity of the CP. In the scenario, the pediatrician pauses the CP (pressing the `Pause` button beside the current activity) and creates a new Therapy Prescription sheet to order the drug. This Therapy Prescription sheet will be listed under all the activities that are associated to the same document template. In the meanwhile, the laboratory has sent the test report and this makes a doctor activate the Test Report Evaluation task. This task is related to the Laboratory Test Report document template; consequently, PRODOC lists all the Laboratory Test Report sheets that are related with the CP instance. At this point, there is only one test report sheet that has been written by the lab. The

⁴ At the present moment, this feature is not implemented in PRODOC.

pediatrician selects it, reads it from within the Data Panel and select the next action accordingly.

The evaluation activity can lead to either the decision of exiting the CP, because the test results exclude the possibility of a meningitis infection (negative case); or of carrying on with the next activity according to the CP schema (positive case). In the negative case, the pediatrician selects the CP Exit event and activates it. At this point, PRODOC proposes to provide a written justification for this decision and then it will terminate the CP instance. Although PRODOC allows users to provide a justification for every deviation from the intended process, it requires a mandatory note only whenever a CP is terminated, since this kind of event is highly critical, i.e., with important consequences for the patient progress. When a CP is terminated, PRODOC comes back to the Treatment activity of the Hospital Stay meta-process related to JD. At this point, the pediatricians can either keep using the whole documentation without the navigation aid of a CP map, create a new CP instance to cope with what turned out is not a regular case of meningitis, or terminate the treatment phase to trigger the administrative tasks of the Discharge phase for JD.

7 Conclusions

In this paper we have illustrated PRODOC, a system conceived: i) To allow designers build the interface of a documental systems starting from the interactions of practitioners with their usual artifacts and not from the data model that makes these interactions meaningful; ii) To support users in browsing and using these electronic artifacts in the light of the work processes they wish to externalize. Our approach is to support clinicians in leveraging process models for what they are intended to be, i.e., as pathways, maps they can rely on to orientate themselves in a wild territory of open choices and clear responsibilities. Since documental practices and working practices are often intertwined and mutually supporting each other – as cooperative work and articulation work usually are – the main tenet of our proposal is that to make documents and processes more integrated can help making practices more aligned to intended reference models of action, a point that at least for more structured models of patient record has been shown of some use Bossen (2006). To gain proof of this tenet, we deployed a prototype in the hospital domain, where patients are the resources being documented on personalized records, and Clinical Pathways are the processes according to which doctors cure patients. That notwithstanding, we believe the concept of PRODOC is applicable to any domain where documents are used as records in order to document events, findings and interventions that regard a specific resource within a practice that actors are supposed to align to a specific protocol, procedure or process.

Irrespectively of where PRODOC can be used⁵, it fosters the externalization of work processes to capture, also in a graphical and visual way, the links between the

⁵ We are planning to deploy PRODOC also in the archeological domain, where the documents under consideration are the excavation records and the resources under documentation are either the stratigraphic units or found artifacts.

procedural aspects of practice and the inputs/outputs that each activity usually consumes/produces. In this way, users can leverage visual active maps to page through data and make apt use of their records. Moreover, if the integration with legacy and organizational information systems is a necessary requirement, PRODOC can be seen as a sort of process-oriented front-end to data that is architecturally “on top of” the legacy system. This integration would then require that the legacy application exposes its data to PRODOC in terms of well defined and bounded “pages” (as in the case of web-based applications). If the underlying system does not have a steady concept of “page”, PRODOC can provide “input forms” to the underlying data management system. Unfortunately, modern document systems have complex DBs inside and usually do not expose them: electronic patient records make no exception. This is for at least two reasons: first, document systems are generally “jealous” of their data, due to justified requirements of data protection and security. Second, the Active Server Pages (ASP) of a web-based document system gather and show data according to the internal state of the application process (i.e., of its workflow or business logic). This means, for instance, that linking process activities to the URLs of the document system’s pages would be useless, even if the navigation system on top of the document system could pass it lots of parameters.

The moral here is that document systems (e.g., EPRs) have their inner workflow, and this is a priori different from, and often irreducible to, the case-specific process that users can adopt in PRODOC. For this reason, and to provide the proof of the PRODOC concept with respect to the tenets illustrated in Section 3, the current prototype embeds an electronic patient record that closely mimics paper-based charts and lets users define and update their process maps in a decoupled manner from document templates. Due to its innovative characteristics, we are now validating PRODOC in a set of informal user sessions from which we gained the first set of findings we report in Section 6. As output of these user sessions, we are considering to improve the interaction design of the user interface, and to extend PRODOC in terms of multi-role and profile management, as well of capabilities of transparent sheet sharing between different concurrent processes.

Acknowledgements

The work presented in this paper has been partially supported by the Italian fund FAR 2008. The authors would like to thank the management and the Internal Medicine and Neonatal Intensive Care Unit personnel of the Manzoni Hospital in Lecco (I) for their kind collaboration. In particular, we would like to acknowledge the invaluable help and courtesy of Dr Bellù.

References

Aigner, W. and S. Miksch (2006): ‘CareVis: Integrated visualization of computerized protocols and temporal patient data’. *Artificial Intelligence in Medicine*, vol. 37, no. 3, pp. 203–218.

- Bardram, J. E. and C. Bossen (2005): 'A web of coordinative artifacts: collaborative work at a hospital ward'. In: *GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. New York, NY, USA, pp. 168–176, ACM Press.
- Berg, M. (1997): 'Problems and promises of the protocol'. *Social Science and Medicine*, vol. 44, no. 8, pp. 1081–8.
- Berg, M. (1999): 'Accumulating and Coordinating: Occasions for Information Technologies in Medical Work'. *Computer Supported Cooperative Work*, vol. 8, no. 4, pp. 373–401.
- Berg, M. and B. Winthereik (2003): *Health Information Management: integrating information and communication technology in healthcare work.*, Chapt. Waiting for Godot. Routledge, London, UK.
- Blaser, R., M. Schnabel, C. Biber, M. Baumlein, O. Heger, M. Beyer, E. Opitz, R. Lenz, and K. Kuhn (2007): 'Improving pathway compliance and clinician performance by using information technology'. *International Journal of Medical Informatics*, vol. 76, pp. 151–156.
- Bossen, C. (2006): 'Representations at work: a national standard for electronic health records'. *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, vol. 37, no. 3, pp. 69–78.
- Boxwala, A., M. Peleg, S. Tu, O. Ogunyemi, Q. Zeng, D. Wang, V. Patel, R. Greenes, and E. Shortliffe (2004): 'GLIF3: a representation format for sharable computer-interpretable clinical practice guidelines'. *Journal of Biomedical Informatics*, vol. 37, no. 3, pp. 147–161.
- Bringay, S., C. Barry, and J. Charlet (2006): 'Annotations: A Functionality to support Cooperation, Coordination and Awareness in the Electronic Medical Record'. In: *COOP'06: Proceedings of the 7th International Conference on the Design of Cooperative Systems*. France, Provence.
- Cabitza, F., M. Sarini, and C. Simone (2007): 'Providing awareness through situated process maps: the hospital care case'. In: *GROUP'07: Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*. New York, NY, USA, pp. 41–50, ACM.
- Cabitza, F., M. Sarini, C. Simone, and M. Telaro (2005): "'When Once Is Not Enough": The role of redundancy in a hospital ward setting'. In: M. Pendergast, K. Schmidt, G. Mark, and M. Ackerman (eds.): *GROUP'05: Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work*. Sanibel Island, Florida, U.S.A., pp. 158–167, ACM Press.
- Cabitza, F., C. Simone, and M. Sarini (2008): 'Knowledge Artifacts as Bridges between Theory and Practice: the Clinical Pathway case.'. In: *KMIA'08: Proceedings of the International Conference on Knowledge Management In Action*. Milan, Italy, IFIP.
- Campbell, H., R. Hotchkiss, N. Bradshaw, and M. Porteous (1998): 'Integrated care pathways'. *British Medical Journal*, vol. 316, no. 10, pp. 133–137.
- Fitzpatrick, G. (2000): 'Understanding the Paper Health Record in Practice: Implications for EHRs'. In: *HIC'2000 Proceedings of Health Informatics Conference, Adelaide, AU, 2000*.
- Fitzpatrick, G. (2004): 'Integrated care and the working record'. *Health Informatics Journal*, vol. 10, no. 4, pp. 291–302.
- Hardstone, G., M. Hartswood, R. Procter, R. Slack, A. Voss, and G. Rees (2004): 'Supporting informality: team working and integrated care records'. In: *CSCW '04: Proceedings of the 2004 ACM conference on CSCW*. New York, NY, USA, pp. 142–151, ACM Press.

- Harper, R. H. R., K. P. A. O'Hara, A. J. Sellen, and D. J. R. Duthie (1997): 'Toward the Paperless Hospital? A Case Study of Document Use by Anaesthetists'. *British Journal of Anaesthesia*, vol. 78, pp. 762–767.
- Hartswood, M., R. Procter, M. Rouncefield, and R. Slack (2003): 'Making a Case in Medical Work: Implications for the Electronic Medical Record'. *Computer Supported Cooperative Work*, vol. 12, pp. 241–266.
- Heeks, R., D. Mundy, and A. Salazar (1999): 'Why health care information systems succeed or fail'. Technical report, Institute for Development Policy and Management (IDPM), Manchester, UK.
- Koppel, R., J. P. Metlay, A. Cohen, B. Abaluck, A. R. Localio, S. E. Kimmel, and B. L. Strom (2005): 'Role of Computerized Physician Order Entry Systems in Facilitating Medication Errors'. *Journal of the American Medical Association*, vol. 293, pp. 1197–1203.
- Lenz, R., R. Blaser, M. Beyer, O. Heger, C. Biber, M. Bäumlein, and M. Schnabel (2007): 'IT support for clinical pathways – Lessons learned'. *International Journal of Medical Informatics*, vol. 76, pp. 397–402.
- Morrison, C. and A. F. Blackwell (2009): 'Observing End-User Customization of Electronic Patient Records'. In: V. P. et al. (ed.): *End-User Development*, Vol. 5435 of *Lecture Notes in Computer Science*. pp. 275–284, Springer-Verlag.
- Mulyar, N., M. Pesic, W. van der Aalst, and M. Peleg (2008): 'Declarative and Procedural Approaches for Modelling Clinical Guidelines: Addressing Flexibility Issues'. *Lecture Notes In Computer Science*, vol. 4928, pp. 335.
- Pedersen, T. B. and C. Jensen (2001): 'Multidimensional database technology'. *Computer*, vol. 34, no. 12, pp. 40–46.
- Quaglini, S., S. Panzarasa, A. Cavallini, G. Micieli, C. Pernice, and M. Stefanelli, 'Smooth Integration of Decision Support into an Existing Electronic Patient Record'. In: *AIME 2005, Proceedings of the 10th Conference on Artificial Intelligence in Medicine, Aberdeen, UK, July 23-27, 2005*. pp. 89–93.
- Schmidt, K. (1997): 'Of maps and scripts: the status of formal constructs in cooperative work'. In: *GROUP'97: Proceedings of the GROUP Conference*. Phoenix Arizona USA, pp. 138–147, ACM Press.
- Shahar, Y., S. Miksch, and P. Johnson (1998): 'The Asgaard project: a task-specific framework for the application and critiquing of time-oriented clinical guidelines'. *Artificial Intelligence In Medicine*, vol. 14, no. 1-2, pp. 29–51.
- Short, M. (1997): 'Charting by exception on a clinical pathway'. *Nursing Management*, vol. 28, no. 8, pp. 45–46.
- Sloan, M. D. and C. S. Guinane (1999): *Analyzing Clinical Care Pathways*. McGraw-Hill Professional.
- Smart, J. F. and M. Roux, 'Medical Knowledge Representation for Medical Report Analysis'. In: *AIME'95, Proceedings of the 5th Conference on Artificial Intelligence in Medicine in Europe, Pavia, Italy, June 25-28, 1995*. pp. 53–64.
- Sutton, D. and J. Fox (2003): 'The syntax and semantics of the PROforma guideline modeling language'. *Journal of the American Medical Informatics Association*, vol. 10, no. 5, pp. 433–443.

Finding the other 5%: Understanding the role of social networking technologies in building personal networks for young adults with cancer

Shawn Ashkanasy¹, Frank Vetere², Hillary Davis³, Graeme Shanks⁴

^{1,2,3,4}University of Melbourne, Parkville, Australia

¹shawna@unimelb.edu.au ²f.vetere@unimelb.edu.au, ³davish@unimelb.edu.au

⁴gshanks@unimelb.edu.au

Abstract: In this paper we explore the ways in which young adults with cancer (aged 17-24) build support networks through computer mediated personal networks. The support networks are influenced by technological affordances and the ongoing experiences of living with the illness and treatment regimes. We report a single, in-depth case study of one young adult with cancer and her use of mobile telephony and web based social networking sites in building support networks. Three important themes emerge from this case. First, in this context computer mediated communications (CMC) are not exclusive to the maintenance of online relationships, but mediate networks of “core”, “significant,” and new ties (primarily online) over time. Second, the social engagement between the subject and members of their social networks is dynamic with different modes of communication predominant at different points in time and different relationships significant at different points in time depending on state of illness, treatment and context. Finally, the interplay between CMC and different ties influence the characteristics of the networks, which is characterized by bridging and segmenting networks.

Introduction

Current literature suggests that people are utilizing communication tools such as the internet to find specialized information and support with people whom they share similar experiences and stories of illness (Burrows *et al.*, 2000; Hardey,

2002). Against this background, we find that young adults with serious or chronic illnesses are going online to find information as well accessing existing and new social relationships through online tools (chatrooms, blogs, forums, social network sites such as Facebook and Myspace). Evidence shows, these technologies are potentially connecting young adults to both existing networks of relationships as well those based on specialised interests or experiences. In a study of online self-help groups for breast cancer, Høybye *et al* (2005, p.216) found that sharing personal accounts of illness (stories) led to transformation and empowerment by offering a mode of action – “Internet based support offer a space for recognition, a social level for sharing knowledge and experience and an individual level for finding ways to live with breast cancer.” Again, Bowker (2008) highlighted in a study of people with disability that the lack of visual and social cues allowed them to meet strangers and converse based on common interests rather than physical, mental or social signifiers. However, the study of people with either a serious or chronic illness has primarily focused on online groups (e.g. www.grouploop.org an online support community for young adults and adolescents living with cancer) who share goals and interests, and promote active participation (Preece & Maloney-Krichmar, 2003).

On the other hand, authors such as Boase *et al* (2006) argue that the use of communication technologies by these users varies according to the changing situation of their condition and health needs although, they are interacting with a circle of core and significant ties that predate their diagnosis and subsequent illness. The author’s define core ties as people with whom the individual has frequent contact, emotional intimacy and availability of network capital. Wellman (2001, p.233) refers to network capital as the availability of resources through ones interpersonal ties that are “widely available, usually specialised, and unevenly distributed among people, ties, and networks.” Significant ties on the other hand, refer to people outside the individual’s core ties. They usually have less frequent contact and are not as closely affiliated and, whilst they are not strangers, their interpersonal importance can fluctuate over time as people access these networks to get help or advice. On a day-to-day basis, the individual may have contact with a variety of core and significant ties through different social networks both face-to-face and computer-mediated.

As such, how people in this situation use communication technologies to access support through different networks and social ties is articulated through their personal networks. *Personal networks* can be described as networks of ties derived from a sample of individuals that enumerates the local social networks (Marsden, 1990), kinds of relationships they contain, and the kinds of resources that flow through different kinds of networks (Wellman, 2007) that shape individual and group action. However, through our research, we also observed that the communication medium itself constrained and influenced the action and social relations between the people and their personal networks (Rice, 1994).

The following paper inquires into the relationship between the communication medium and the support needs of the individual and how this influences the individual's personal support network (which is a subset of their overall network). This is explored through empirical evidence from a recently completed study of a group of young adults being treated for a range of cancers and at different stages of treatment. In particular, we focus on the life of one of the participants, Bianca, and her use of computer mediated communication (CMC) in accessing support and information. Firstly, we explore her cancer experience and how it affected her support and informational needs. We elaborate upon the circumstances of her illness and how it shaped to some degree her communication with different members of her network. Secondly, we discuss the interdependency between different social ties and communication through CMC. We develop this further by providing different examples of her support needs and how this influenced the way she communicated and to whom, including her use of technology and how this also influenced her personal network. Finally, in the discussion, we draw out the relationship between the communication medium and the support needs of the individual through Wellman's theory of *networked individualism* (Wellman, 2001). By placing this socio-technical relationship within a social network discourse, we argue that Wellman's theory provides new insight into the factors that shape the use of collaborative technologies and, in particular social support through social networking technologies.

We begin by elaborating upon current theoretical work in CSCW that addresses personal networks and collaboration. This is followed by the research design and the case study. Subsequently the case is analysed using networked individualism, including implications for the study of personal support networks and CSCW.

CSCW and personal networks

The use of personal networks in CSCW in both workplace studies, as well as those that address social support are scant. Nardi *et al* (2000) provides, perhaps, the best example of a workplace study that uses personal networks. They found that individuals within the workplace managed their own personal networks of work contacts with whom they collaborate with over short or long periods of time. These "intentional networks" as the author suggest, are not well understood through existing explanatory models of workplace collaboration – "our study documents the wider, less predictable, set of social relationships in which workers are implicated" (Nardi *et al.*, 2000). They argue the structure of intentional networks is not likely to be based on common experience, unlike existing notions of community. Rather, they are oriented around the individual and not necessarily based on an ongoing commitment; they are widely distributed rather than based on a fixed location; they are more flexible, yet organized in deliberate

ways by the individual; they are affected by temporal patterning, whereby the network can transform over time.

As such, intentional networks link the individual's immediate work community with a much more distributed and temporal one. The two social worlds exist simultaneously, while affording different social opportunities and access to resources. At any point in time, the individual may be engaged in any number of social networks with specific properties, members, types of support and forms of communication. Furthermore, the networks are not mutually exclusive insofar as they may share members, whereby the communication medium itself may connect more than one network, while other communications mediums may exclude or segment the network.

Conversely, the tendency in CSCW has been to investigate local groups of potentially knowable members from either an organizational setting or local geographically situated communities of citizens. An early example can be seen in the work of Mynatt *et al* (1998) who studied three multi-user domains or MUDS, which connected members of a relatively small network of local ties within an organizational setting. They argued that certain communication media were more apt to provide users with a sense of virtual place than others by virtue of their persistent, flexible and multi-user environments. These "networked communities" (Mynatt *et al.*, 1998, p.123) are limited at least in theory to understanding community as local forms of social groups that are bounded, small and characterized by dense social relationships. Similarly, more recent examples of work in CSCW that address community building, participation, and information sharing can be seen in Kavanaugh *et al* (2007) and Munkvold and Ellingsen (2007) respectively. Whilst both these authors address the interplay between technology and interpersonal relationships, the emphasis on group level interaction neglects the broader social and technological setting in which these relationships are embedded. This is less a criticism and more an acknowledgment of the limitations of this approach to understanding the dynamic structure of personal networks and how people use them to socialise and find support.

Research Method

The following account uses an interpretivist approach to draw out themes and outcomes from a case study of one participant's interviews and field observations. The unique characteristic of young adults with cancer (YAWC) makes studying this group often difficult in the sense that their lives are to an extent unpredictable and volatile. There is strong evidence that cancer amongst young adults ranging between 18-24 years of age is one of the most complicated conditions. This is due to a combination of rare and invidious cancers that affect young adults as well as many psychosocial aspects associated with this stage of life (Bleyer, 2002). These psychosocial concerns are social and sexual development, identity and self image,

work and education, goal achievement, peer pressure, intimacy, fertility and reproduction, and the need for independence and autonomy (Bleyer & Barr, 2006; Bleyer, 2002; Ettinger & Heiney, 1993; Thomas *et al.*, 2006, Woodgate, 1998). Hence, it is generally acknowledged by oncologists and psychologists working with this group that the psychosocial aspects of cancer are not only unique amongst the age group, but they are also broader in their scope than in the rest of the population (Bleyer, 2002). This would also account for the relatively limited number of studies that address their communication needs and practices. As such, we argue that a single case study of this nature fulfils Yin's (2002, p.40) rationale as "revealing," whereby the situation under investigation has been previously inaccessible to systematic observation (Yin, 2002) and the description of information alone is revelatory.

The findings examined in this paper are the result of empirical data collected from a recently completed study of YAWC. This paper refers to a single in-depth case of one patient, which is part of a larger body of research. The unit of analysis for the study is one particular journey through diagnosis and relapse over a two-year period. The data presented in this paper was collected over a three-month period and were taken from three sources. Firstly, three semi-structured interviews were conducted starting in early June 2008; sections of these will be presented in the following. Secondly, observations were made from a purpose-built web-based self-help site for YAWC, called *MyTrac*. Observations from the participant's use of the site were taken both during and after the completion of the study, which also provided reference points for the interviews. Thirdly, electronic logs were generated from messages sent through *Twitter*, which is a mobile phone to web-based micro-blogging service. *Twitter* allows the user to provide status updates from their phone or computer to their own *Twitter* page as well as other users who are 'following' them. The system distributes these messages to a potentially infinite network of connected followers.¹

In the context of this paper we distinguish between the roles of *MyTrac* in relation to other more 'familiar' modes of web-based social communication, such as social networking sites like Facebook and Myspace. Given that we distinguish our research aims from studies that focus on POC and online communities, we also apply this distinction to the participant's use of *MyTrac* in relation to, for example, Facebook (i.e. the former being an example of a POC). As is borne out in the data the participant's use of *MyTrac* is clearly different to her use of Facebook and, similarly, her use of *Twitter* in relation to her other mobile communication practices. As such, much of the data presented focuses on her use of Facebook and other tools such as Instant messenger (IM), which she used actively before and during the study. References to *MyTrac* and *Twitter* provide supplementary evidence of her personal networking across groups and mediums.

¹ *Twitter* has since ceased providing the update to mobile phone service outside America, India and England.

Data collection and analysis

A qualitative approach was used in order to understand the particular context and setting in which CMC are used. We selected to focus on Bianca because she was the most “appropriate participant” (i.e. those who can best inform the study). In other words, qualitative inquiry points to depth rather than breadth: “In qualitative research we are not interested in an “on average” view of a patient population. We want to gain an in depth understanding of the experience of particular individuals or groups” (Greenhalgh & Taylor, 1997, p.741). Themes presented in this paper are the result of an iterative process whereby larger more boarder themes were compared and contrasted to express patterns of Bianca’s communication between the participant and her personal network. In the discussion we draw upon networked individualism, which is used to both describe the participant’s personal network as well as account for the interplay between CMC and their communication behaviour.

Case – Bianca and finding the other 5%

Bianca was 20 when she was diagnosed with Hodgkin’s Lymphoma in 2006 and relapsed at her six-month check up in 2007. Hodgkin’s Lymphoma as her Oncologist informed her is one of the most treatable cancers, with a 95% success rate. With another round of chemotherapy, Bianca’s treatment succeeded in removing almost all of the cancer except one near her heart. The stem cell transplant that followed also failed to remove the cancer. Running out of options they attacked the cancer with radiation, focusing on her mantle (lower rib cage to her chin). Unfortunately, the treatment did not work and in fact the cancer spread to her sternum, her lung bases and abdomen. She was, as she described in the interviews, part of the 5% of patients with Hodgkin’s Lymphoma that do not respond to treatment. At the start of the study Bianca was waiting for a bone marrow match for another transplant and she was on a waiting list for a clinical trial. To suppress the growth of the cancer she took high doses of steroids, but the drugs often caused nausea as well as lowering immune system. This meant that leaving home or for that matter physical interaction was fraught with risk of either vomiting or infection.

Cancer and indeed her support and informational needs changed when she relapsed. Bianca recalls:

[Int.1] If I see people... I know what it was like for me the first time around, I had my friends and family and that’s great, but speaking to people who have gone through it. It was fine the first time round because you keep getting told about these odds that 95% your cured, um, well, your in remission and then five years later your cured, but the chances of relapse are really small, so you go along deluded, I guess, in some way that you’re in that 95% category. And then to find out that you’ve relapsed and that you have relapsed so quickly and that it has come back so aggressively, um, is terrifying and then you’ve only got that 5% to work with. Where

are these 5% and then you get thinking and you try find these people and that sort of what spurred me on more afterwards, people knew that I had it, obviously I don't say hello I've got cancer but it's quite obvious in my page (Facebook page) if people read between the lines, like "how's your chemo going", well chemo's associated with cancer, um but yeah, that's what really made me look beyond, I guess to find people going though the same thing, because I guess it's definitely different the second time round.

The focus of Bianca's story is what happened after she relapsed, the changes that occurred to her network of supportive ties, where and what she communicated and how this was supported by CMC. Much of Bianca's support prior and post relapse was from family, family friends and close friends with whom she had regular contact. The marked differences that relapse brought was a gradual but sustained presence of online contacts that she had made through various cancer related groups:

[Int.2] When I was first diagnosed I was not using any of the networking; I wasn't on Myspace or any of that. I think it was part way through the first time that I joined Myspace through a friend and then when I was in the (name omitted) I joined Facebook and then they took over.

As Bianca suggests in this extract, Facebook was a considerable part of how she found and made connections with other cancer patients in different locations, both locally and globally. The affect, as such, largely contributed to her ability to cope with the uncertainty of relapse and the questions that arose out of that experience:

[Int.1] I think it has been a huge positive being able to connect to someone that has had it or has had it relapse because just knowing you're not alone.

These weak ties played an important role when stronger ties were unable to provide information. To reiterate Bianca's comment regarding finding other young adults with cancer when she relapsed:

[Int.1] I know what it was like for me the first time around, I had my friends and family and that's great, but speaking to people who have gone through it

As such, weak ties link people with different social characteristics and knowledge, who are more apt to provide new information (Wellman and Gulia, 1999). In this way, online social networking tools can provide opportunities for previously disconnected people to make direct contact:

[Int.1] How I got into Myspace and through Myspace I got into Facebook and I've met people who are actually in the same boat as me that have not responded to treatment with the same cancer and similar aging and you can compare with them, "what trials is he being offered"? "What trials are working"? What have you heard about?

However, through her use of this social networking site she has also integrated other friendships and associations into this digital domain. When we interviewed Bianca, she said that she had approximately 230 'friends' in Facebook and of those she maintained regular communication with around 50:

[Int.3] Like I said, some of them are from primary school and high school that, um, that you knew, but you were not heaps close with, um, so I don't mind if they read about it that's me, I can't change what I'm going through, um, but yeah, there are some people that I do

communicate with a lot more whether they are patients or closer friends, um, others are there and you do know them, but you don't communicate, so probably about 50, it's hard to gauge

Facebook in this context operated like a network exchange linking Bianca's relatively vast network of social ties within a singular digital place. However, in the sense that Bianca's 'friends' were linked by virtue of a common thread, i.e., Bianca, the network was not determined by its locality. Rather, Bianca's use of Facebook supports Wellman's contention that "each person operates a separate personal community, and switches rapidly among multiple sub-networks" (Wellman, 2003). In this way Facebook was conducive to her way of receiving support and information, as well as giving it:

[Int.3] I'm not part of a big group though, a lot of my friends are more individual contact; I never really did fit in with the whole group thing.

We also found that MyTrac was less a single community, and rather a collection of individuals who maintained their own networks of which MyTrac was apart. Bianca revealed that her interactions with the other MyTrac members was predominantly one-to-one and therefore more oriented towards her reaching out to communicate or being directly contacted herself:

[Int.2] MyTrac is more comments on people's pages, it's weird, I guess it's because we have nothing to hide. You already know it it's written on your page

In this way, MyTrac and Facebook afforded opportunities for both focused interaction between Bianca and her network of core and significant ties as well as "crosscutting" (Wellman, 2001, p.234) that linked and integrated her social milieus.

From this description of Bianca's communication practices, it is possible to see that different technologies played different social roles and that this was mediated by her support and informational needs. In addition, the motivation to use a specific communication medium was influenced by the type of support available through particular ties linked to specific modes of communication (Haythornthwaite, 2002). However, what is particularly interesting here is how different mediums influenced, on the one hand, the kinds of ties prevalent, and on the other hand, the kinds of support she received. To best illustrate this we identified two key aspects of her communication practices, namely *bridging* and *segmenting*. Whilst bridging was clearly evident in the notion of "crosscutting" networks, segmenting emerged from further data analysis and more clearly characterized the way different communication technologies demarcated specific relational ties.

Bridging networks

Insofar as Boase and Wellman (2006) argue that crosscutting ties links and integrates social groups, increases societal connectivity, we can see that Bianca's use of Facebook facilitated people coming together, albeit it was oriented by specific temporal events. As such, Bianca's communication through Facebook

fostered the folding in of this online network into broader communities, what we have called bridging. The following example of bridging took place when Bianca received results from a scan that her cancer had gone into remission (although remission is only properly applied after two years of no symptoms). Although what makes this particular instance unusual was the timing of her disclosure, as Bianca chose to inform the MyTrac community first via Twitter on her mobile phone, in which case the members of MyTrac replied virtually instantaneously, from multiple locations:

Bianca: Just arrived at the (name omitted) 4 my appt wit the Prof.. I've got an upset tummy, nerves. Hope news isn't 2 bad.

Fiona: Good luck Bianca!! Im waiting around at RMH for a Doc and a bed. Still not sure what's going on.

Thomas: Good luck 2 the both of u, I've gone thru 8 years of operations n pain every year, just think positive n u will get thru it. I will say a little prayer 4 u

Bianca: Really skeptical. The last scan came back all clear. WHAT! All I've had is 1 lot of steroids. This cancer is messing wit me. Another scan next wk

Bianca: all questioning it, Prof incl.. it's odd

Thomas: Take out the champagne n celebrate!! That's good news that the scan is clear isn't it?

Bianca: Test result odd. Having further tests next week. Nothing definite.

Thomas: Well I hope the further tests show up clear 2 then =). Good night n sweet dreams all.

Bianca suggested in the interview that she needed to “send something” and the members of MyTrac were new contacts that had an understanding of what she was going through and furthermore it was a smaller network and therefore easier to manage peoples expectations:

[Int.2] I didn't tell everyone straight away because I didn't want to get hopes up and I didn't know what was happening, but I wanted to send something and then I thought, I'll send to Twitter, but then it cut out part of my message, so then I had to rephrase it because everyone thought I was in remission. "No I'm not!" So I had to resend it, but it was good and it was good to have that feedback, I wasn't expecting it

She then reflected on the response, and reconstituted the information for a different, wider community on Facebook through her ‘status bar’:

[Int.3] I did tell people and I have had it on Facebook that it's all clear at the moment, AT THE MOMENT! So people do know that the last one was clear, but I just didn't go through my whole phone like I used to. Like say, if I got a good result, I'd go send to that person and that person that person... and I'd just go through the majority, but because there were still so many questions and it's only one scan and I'm getting scanned again next month...

However, the affect of this was quite different, Bianca notes that her friends on Facebook understood the significance of her good news differently:

[Int. 3] Yeah, like some did (understand it), but they got excited because at least finally it was some good news for me, I think that was the other reason, they were just really glad that finally something has gone my way because so much has gone against me. To have something positive happen for once (was good)

In this case, the audience was not only much larger than on MyTrac, but she did not differentiate between network members, and as she later stated – [Int.3] *If*

they don't want to read it, don't go on my page. The medium, in this instance reduced the boundaries between networks and linked network members through their common interest in Bianca's wellbeing. This also supports Granovetter's (1973, p. 1376) claim that "the more local bridges in a community and the greater their degree [of density], the more cohesive the community and the more capable of acting in concert", although in this case the medium was a significant player in enabling her community to act in concert.

Segmenting Networks

In the sense that Bianca used Facebook to momentarily bridge her networks, we also found that the communication medium as well as other temporal factors segmented them, whereby demarcating the strength of relational ties. This was evident from our interviews where illness and treatment played an important role in determining the mode of communication as well as people with whom she communicated. She remarked that while she was in hospital for chemotherapy and radiation she was often too sick to use her computer to stay in touch with people:

[Int.3] It's whether you've got the energy to and sometimes you are just sitting there and I just want to go to sleep or I just can't be bothered doing this, or you just don't have the heart for this. It's not so much bringing it all back, you don't have the energy, it's almost like it's too hard and you say it can wait, it will be there when I check it next, so you don't go on. When I was in the (name omitted) I didn't check my internet unless I was really bored and I felt alright, because if you're being sick and your tired, or you've got visitors or something it's really not ideal.

In this instance, not only did less significant ties recede to the background, but the mode of communication was also minimized and therefore the type of people she communicated with. In this case the medium not only reflected the relational bond between Bianca and her communication partners, but also momentarily reduced the size of her network. This was revealed in two separate conversations with Bianca. The first concerned her father and the importance of using her mobile phone to bridge the gap between times of absence:

[Int.1] When I was in for the stem cells [...] Our (referring to her father) form of communication would be on the phone, we'd speak everyday, but it's still not the same because he's not there and as much as he'd want to be there for you he can't because he's got to work

In contrast to the need to sustain contact with her father, her commitment to less intimate ties, such as other patients she met online, became momentarily less important:

[Int.1] I used to email updates to people but I sort of just... people were relying on that and they were just waiting for an email and they would contact me and I thought, "it's so impersonal" and then I gave up on that and with all the radiation (it became too difficult)

What this suggests is that the communication medium can play an important role in not only affirming the significance of a particular tie, but also the medium

can segregate networks and enforce the strength of relational ties and the members' place in the network. This was also evidenced where Bianca used more than one communication modality to support communication with core and more significant ties. On a daily basis, Bianca used a combination of mobile phone and instant messaging software (i.e. MSN) to communicate with particular network members. In the follow extract Bianca talks about how she communicates across channels to her best friend and the role these mediums played:

[Int.3] Um, say I talk to someone though text or I see them in person or whatever. I went and saw (name omitted) yesterday and she sort of said I will be online later, so if you are online we will chat, so sometimes you will continue conversations, or if you left a comment on someone's page and then they logged onto MSN and then you would start following on from that, um, it's just a way of communicating. After having it for so long, you don't even think about it, you're in that generation that you've grown up with it, like I've had MSN since year 8, so you've got a fair few contacts on that

However, moving between mediums was not simply about physical or temporal constraints, but a mixture of finding the right medium to support the emotional and informational content of the relationship. For example, of the young adults with cancer she had met online she established three significant relationships with women who have also not responded to treatment. The most intimate of these lives in the UK and they maintain regular communication through phone messaging, MSN and Facebook. Also, through the study and her use of MyTrac, Bianca became close with one of the other participants, Louise, who also had Hodgkin's Lymphoma, but had not relapsed. Bianca's interaction with Louise during the study was initially through MyTrac, but as their friendship became more intimate they communicated through MSN and SMS as well some interaction through Facebook:

[Int.3] I guess initial common ground was Hodgkin's lymphoma, same thing, and then we added each other to MSN. Sometimes we don't even talk a whole lot, or sometimes we talk about everyday things or stupid things that have happen or if we're watching TV while we are talking to each other, we'll talk about that, there are different wink things you can do that are quite amusing to just annoy the crap out of each other. Sending them and always "ding ding", and then all of a sudden it comes up and all of these people are hammering (referring to the emoticon) and doing these things on the screen. I don't know, it's just amusing, but we get along really well. Like if she is online we pretty much speak, um, I guess most days if not every couple of days

Insofar as the medium needs to support the informational and supportive content of the relationship, the content itself needs to reflect the multiplicity of the relationship. This was also embedded in the type of medium used to communicate, but also how she demarcated someone's relational significance:

[Int.3] I take it as it comes, if someone is interested, whether it's online or...It is a lot more personal when you speak to someone in person or on the phone, but there are some people where your form of communication with them is online that's the basis... I guess that's why

they're not I guess your close friends because that's the only form of communication you have is the online stuff

This was again reinforced when we asked Bianca about another member of MyTrac. What was interesting about Bianca response was that she suggested there was equivalence between the nature of her relationship with this young person and the mode of communication:

[Int.3] But yeah, I sort of spoke to him, but again it's just someone to speak to in that sense it's not, um, it hasn't gone any further like with MSN or any other networks as a form of communication.

Bianca's use of one-to-one (email, mobile phone, MSN), one-to-many (Twitter), and many-to-many (Facebook, MyTrac²) forms of communication enabled her to build and maintain her personal network that was not only distributed over and between physical and virtual space, but was amenable to the many temporal and physical constraints of her illness. Furthermore, the medium was an important agent in changing the characteristics of her personal support network as well as reinforcing specific relational ties.

Discussion

From this case study of Bianca's personal support network, we derived three key findings that describe the socio-technical relationship between communication technology, support and informational needs, and specific social ties:

- i) In this context computer mediated communications are not exclusive to the maintenance of online relationships, but mediate networks of "core", "significant," and new ties (primarily online) over time.
- ii) The social engagement between the subject and members of their social networks is dynamic with different modes of communication predominant at different points in time and different relationships significant at different points in time depending on state of illness, treatment and context.
- iii) The interplay between CMC and different ties influence the characteristics of the network, which is characterized by bridging and segmenting networks.

In the following, we analyse these findings using networked individualism (Wellman, 2001).

² Both MyTrac and Facebook encompass individual and group communication modalities. Our categorization schema is based on Bianca and may change for other users.

Networked individualism and personal networks

With the increased mobility of both people and culture brought about by the proliferation of new technology driven communication— mobile phones, internet – people have access not only to diverse forms of information, but also new relationships and opportunities for sociability and support that they might not otherwise find in their local support networks of core friends and family members. Authors such as Wellman (1999, 2001) and Castells (2001) have argued that society and people more generally have moved away from centralized and geographically located communities towards “social networks organized by shared interests rather than by shared locality” (Wellman, 1996, p.5). Networked individualism (Wellman, 2001, p.248) suggests that each person maintains their own “personal community” of multiple, thinly-connected, and partial communities, which they can switch between kin, neighbours, friends colleagues and other specialized relationships (Wellman, 2001, 2003). This social arrangement provides the basis through which individuals seek and maintain relationships that afford amongst other things support, information, social identity and belonging (Wellman and Gulia, 2001; Wellman, 2003). Rather than fitting into the same group as those around them, people potentially accesses a vast social network that spans and intersects local and global localities, mediated and face-to-face. Boase and Wellman (2006, p.720) have stated the core characteristics of networked individualism as:

- i) Relationships are both local and long distance.
- ii) Personal networks are sparsely knit but include densely knit groups.
- iii) Relationships are more easily formed and abandoned.

Firstly, the data presented provided evidence that Bianca’s personal support network is not limited to any single group or setting. Rather her relationships spill over into multiple, partial communities that comprise, as Boase and Wellman suggest, *both local and long distance relationships*. Furthermore, these were supported through a suite of communication mediums that afforded different types of relationships and interactions. For example, those relationships with a high social / emotional commitment were often supported through the use of multiple technologies, than is indicated by just using email for example (Davis, Vetere & Ashkanasy, 2008). This, as Wellman argues, refocus attention from fixed groups to “active networking” in order to find support, sociability and identity (Wellman, 2001, p.234).

Secondly, Bianca’s use of one-to-one, one-to-many, and many-to-many modes of communication changed in response to external factors linked to illness and treatment, as well as internal social dynamics between herself and her relational ties. As such, her communication practices support Boase and Wellman’s assertion that people communicate between *many sparsely knit networks and densely bound groups*. However, it is important to recognize that the medium

itself delimits the potential range of the audience as well as demarcating roles, drawing different relationships to the foreground, while necessarily allowing for others to fall to the background. Licoppe and Smoreda (2005, p.317) argue that different communication tools “provide new resources to negotiate individual timetables and social exchanges, making it possible to adjust roles, hierarchies and forms of power in relational economies.” Although we would add the division of roles and relational ties was more clearly evidenced, where the technical constraints reinforced tight boundaries and individual / group sentiments (Wellman, 1996), for example where the audience was smaller and either one-to-one or one-to-many.

Thirdly, whilst we agree with Boase and Wellman *that relationships are more easily formed and abandoned* online than face-to-face, we feel that it was subtler than what they suggest. It was unusual for Bianca to entirely sever a tie, rather weaker ties, or those that provided specific types of information or support, fluctuated in their presence. We feel that the notion of bridging and segmenting describes the rhythms of her relationships, whereby the interplay between events, timing, and her support and informational needs mediated the significance of different relationships; although she was communicating regularly with her more core ties (i.e. her parents and close friends). Furthermore, the immediacy of the communication technology to both send and receive responses was an important agent in this process. For example, the asynchronous nature of her messages on Facebook gave her control over the timing of her disclosures and the audience to whom she communicated. This also allowed non-active members to respond, which has the potential to change their relational significance to the individual and others in the network. Wellman has discussed this in terms of a general reciprocity, whereby comments made in a publicly accessible domain is potentially seen by the entire group and moves to solidify group sentiment and foster positive reward of its members (Wellman & Gulia, 1999, p.176). Also the nature of the disclosure, whether it is to a single person or the group can connect previously disconnected people through their mutual interest in the individual, turning an “indirect tie into a direct tie” (Wellman, 1996, p.6).

This movement between different types of ties and support has also been explored in the work of Foth and Hearn (2007), who studied the communication and social interaction between residents of three inner-city student apartments in Australia. They proposed the concept of “communicative ecologies” (Foth and Hearn, 2007, p.751) to capture the interplay between online and offline, global and local as well as collective and networked social communication. Their study found that communication between residents of these apartments moved between small groups and networks depending upon purpose and context. Foth and Hearn’s (2007, p. 768) concluded on the one hand that “individuals in networks give rise to emergent collective behavior” and on the other hand that peer-to-peer communication encapsulated this idea more so than the use of public forums. The

latter findings supports our argument that one-to-one or one-to-many modes of communication more clearly reinforced tight boundaries and individual / group sentiments. This is significant because it raises questions about the use of, for example, forums in online peer support sites that attempt to encourage group collaboration/sharing. Insofar as different mediums enable different levels of social interaction, they also signify different levels of sociability. Depending upon the context in which they are used, their use, either by an individual or a group can demarcate social and relational boundaries. As such, the relational significance of communication technologies is an important aspect of collaborative communication.

Limitations

This paper is limited by the reliance upon one case study; as such the findings are idiosyncratic and highly dependent upon the experience and interpretations of one person, as well those of the researchers. Other social factors such as gender, socio-economic background have also been neglected, however they require further analysis; a fertile ground for future papers, or indeed studies.

Conclusion

Networked individualism provides a valuable tool to unpack personal networks and the independency between people, technology and relationships albeit it also depends upon the socio-relational context in which they function. Bianca's story provides clear evidence that the internet not only supported interaction with online groups, but incorporated an array of people from all aspects of her life; from core intimate ties which she regularly interacted with physically and virtually, to a larger set of significant ties that fluctuated in their presence in her life, to many more weaker ties that she has met through online groups as well as on Facebook. While the study of group social dynamics and behaviors are an important area of research, personal networks and networked individualism attempt to explore the factors that constrain and influence individuals and the various ways they access resources (both social and material) through different social ties. As such, this tension between the actions of individuals and the characteristics of the networks, including the communication mediums that support them, is an important aspect of personal networks that influences the way support is understood in this context.

Finally, the study of personal networks in CSCW and in particular the use of CMC by people living with illness and how they access support and information must consider a number of factors. On the one hand, it is important to consider the temporal aspects of illness and how they influence the support and information needs of the individual, while on the other hand, the motivation to

use a specific communication medium is also influenced by the type of support available through particular ties linked to specific modes of communication. Furthermore, researchers should bear in mind the way communication technologies demarcate relational ties within personal networks, whereby influencing the type of ties available and support prevalent over time.

Acknowledgments

We would like to thank all of the participants of the study for their generosity of time and in particular, Bianca, for showing us the capacity of the human spirit to survive despite all adversity. We would also like to acknowledge our funding partner, Cancer Australia, and research partners David Thomas, Penny Schofield, Kate Thompson, Gavin Dyson, and Giselle Withers, as well as Telstra, for providing equipment and technical support; without which the study would not have been possible.

References

- Boase, J., Horrigan, B. J., Wellman, B. & Rainie, L. (2006) *Pew Report: The Strength of Internet Ties*. Washington, Pew Internet and American Life Project.
- Boase, J. & Wellman, B. (2006) Personal Relationships: On and Off the Internet. In Perlman, D. & Vangelisti, L. A. (Eds.) *Handbook of Personal Relations*. Cambridge, Cambridge University Press.
- Bleyer, W. A. (2002) Cancer in older adolescents and young adults: Epidemiology, diagnosis, treatment, survival, and importance of clinical trials. *Medical and Pediatric Oncology*, 38, 1-10.
- Bleyer, A. & Barr, R. D. (2006) Highlights and Challenges. IN Bleyer, A., M, O. L., Barr, R. D. & Ries, L. (Eds.) *Cancer Epidemiology in Older Adolescents and Young Adults 15 to 29 Years of Age, Including SEER Incidence and Survival: 1975-2000*. Bethesda, MD, National Cancer Institute.
- Bowker, N. (2008) Participating in the World (Wide Web): Social Connections for People with Disabilities IN Holland, S. (Ed.) *Remote Relationships in a Small World*. New York, Peter Lang.
- Castells, M. (2001) *The Internet galaxy: reflections on the Internet, business, and society*, New York; Oxford, Oxford University Press.
- Davis, H., Vetere, F. & Ashkanasy, S. (2008) Towards Social Connection for Young People with Cancer. *Proceedings of the 19th Australasian conference on Computer-Human Interaction: Designing for Habitat & Habitus*. Cairns, Australia, ACM.
- Ettinger, R. S. & Heiney, S. P. (1993) Cancer in adolescents and young adults. Psychosocial concerns, coping strategies, and interventions. *Cancer*, 71, 3276-80.
- Foth, M. & Hearn, G. (2007) Networked Individualism of Urban Residents: Discovering the Communicative Ecology in Inner-City Apartment Buildings. *Information, Communication & Society*, 10, 749 - 772.
- Granovetter, M. S. (1973) The Strength of Weak Ties. *American Journal of Sociology*, 78, 1360.

- Greenhalgh, T. & Taylor, R. (1997) *Papers that go beyond numbers (qualitative research)*. *BMJ*, 315, 740-3.
- Haythornthwaite, C. (2002) Strong, Weak, and Latent Ties and the Impact of New Media. *The Information Society*, 18, 385-401.
- Høybye, M. T., Johansen, C. & Tjørnhøj-Thomsen, T. (2005) Online Interaction. Effects of Storytelling in an Internet Breast Cancer Support Group. *Psycho-Oncology*, 14, 211-220.
- Kavanaugh, A., Zin, T., Rosson, M., Carroll, J., Schmitz, J. & Kim, B. (2007) Local Groups Online: Political Learning and Participation. *Computer Supported Cooperative Work (CSCW)*, 16, 375-395.
- Licoppe, C. & Smoreda, Z. (2005) Are social networks technologically embedded?: How networks are changing today with changes in communication technology. *Social Networks*, 27, 317-335.
- Marsden, P. V. (1990) Network Data and Measurement. *Annual Review of Sociology*, 16, 435-463.
- Mynatt, E. D., O'day, V. L., Adler, A. & Ito, M. (1998) Network Communities: Something Old, Something New, Something Borrowed. *Computer Supported Cooperative Work (CSCW)*, 7, 123-156.
- Munkvold, G. & Ellingsen, G. (2007) Common Information Spaces Along the Illness Trajectories of Chronic Patients. *Ecscw 2007*.
- Nardi, B. A., Whittaker, S. & Schwarz, H. (2000) It's Not What You Know It's Who You Know. *First Monday [online]*.
- Preece, J. & Maloney-Krichmar, D. (2003) Online Communities: Focusing on Sociability and Usability IN Jacko, J. & Sears, A. A. (Eds.) *Handbook of Human-Computer Interaction*. Mahwah, NJ, Lawrence Erlbaum Associates Inc. Publishers.
- Rice, R. E. (1994) Network Analysis and Computer-Mediated Communication Systems in Wasserman, S. & Galaskiewicz, J. (Eds.) *Advances in social network analysis: research in the social and behavioral sciences* Thousand Oaks, California, Sage Publications.
- Thomas, D. M., Seymour, J. F., O'brien, T., Sawyer, S. M. & Ashley, D. M. (2006) Adolescent and young adult cancer: a revolution in evolution? *Internal Medicine Journal*, 36, 302-307.
- Walsham, G. (1995) Interpretive case studies in IS research: nature and method. *European Journal of Information Systems*, 4, 8.
- Wellman, B., Quan-Haase, A., Boase, J., Chen, W., Hampton, K., Isabel Isla, D. D. & Miyata, K. (2003) The Social Affordance of the Internet for Networked Individualism *Journal of Computer Mediated Communication*.
- Wellman, B. (2001) Physical Place and Cyberplace: The Rise of Personalized Networking. *International Journal of Urban and Regional Research*, 25, 227-252.
- Wellman, B. & Gulia, M. (1999) Virtual Community as Community: Net Surfers Don't Ride Alone. In Smith, A. M. & Kollock, P. (Eds.) *Communities in Cyberspace*. London, Routledge.
- Wellman, B. (1996) For a Social Network Analysis Of Computer Networks: A Sociological Perspective On Collaborative Work And Virtual Community. *Proceedings of the 1996 ACM SIGCPR/SIGMIS Conference on Computer Personnel Research*. Denver, Colorado, United States, ACM.
- Wellman, B. (2007) The Network Is Personal: Introduction to a Special Issue of Social Networks. *Social Networks*, 29, 349-356.
- Woodgate, R. L. (1998) Adolescents' perspectives of chronic illness: "It's hard". *Journal of Pediatric Nursing*, 13, 210-223.

Yin, R. K. (2002) *Case Study Research: Design and Methods*, Thousand Oaks, California, Sage Publications.

Trust and Social Capital: Revisiting an Offshoring Failure Story of a Small German Software Company

Alexander Boden¹, Bernhard Nett¹, Volker Wulf^{1,2}

¹Institute for Information Systems and New Media, University of Siegen

²Fraunhofer Institute for Applied Information Technology, St. Augustin
{ alexander.boden | bernhard.nett | volker.wulf }@uni-siegen.de

Abstract. While work organization and social capital are known to be important factors for offshoring success, there is little empirical evidence on how these aspects evolve in the course of offshoring projects. In the literature, trust has been discussed as a personal disposition to abstain from control in a given situation, and was found to remain surprisingly stable in some cases. By analyzing the relation between control and trust in the course of a failed offshoring project, we want to add to the discussion on social capital as a factor for successful offshoring. The results of our long-term ethnographic study are somewhat paradox: in our case, ongoing conflicts motivated attempts to strengthen control, although personal trust and social capital remained strong. Despite the fact that the confidence of the partners in their offshoring project was weakened over time, the trust among the partners prevailed. However, social capital was not only unable to save the offshoring project—it also seemed to hinder the conflict resolution in some regards. Therefore, we argue that while social capital is an important factor, it should not be regarded as a context-free asset, but rather (in Bourdieu's perspective) as a risky investment.

Introduction

With ongoing globalization, offshore software development has become quite common. For instance, consulting agencies promote Global Software Engineering

(GSE) as a means to reduce costs and as a driver for process improvements in case activities are reengineered and streamlined as part of the move. However, while wage differences may offer options to reduce costs, the spatial, temporal and cultural issues in globally distributed cooperative work are still challenges and need to be better understood (Cataldo et al. 2006; Gutwin et al. 2004; Herbsleb et al. 2000).

Tackling these issues, it is often argued that GSE needs formalization of processes and a high level of social capital to be successful (Levina & Vaast 2008). Features such as formalization and social capital accumulation may—to a certain degree—be influenced when establishing the offshore cooperation. However, there has been little empirical evidence on this topic for later stages of offshore cooperation (King & Torkzadeth 2008). This is even more astonishing as these factors are likely to affect flexibility which is regarded as a major demand for software development in general (and especially for small enterprises). Therefore, we need to learn more about how the relationship among clients and vendors evolves within offshoring projects and which factors contribute to or oppose efficient cooperation (Fisher et al. 2008).

CSCW has a long tradition of researching problems of distributed cooperation. For example, CSCW studies have expounded the importance of awareness, tool appropriation, self-organization, behavior, interaction and communication in different kinds of work groups by means of ethnographic studies. However, there are very few in-depth studies which look at the particularities of cooperative work in off-shored software projects—specifically when small companies are involved.

In order to add to the understanding of offshoring, we conducted a long-term and in-depth ethnographic case study in a small German company between 2006 and 2008. During this time, the company developed software in an international team of German and Russian developers. In the end, the cooperation was terminated due to ongoing problems. By revisiting this failure case and the related conflicts over a longer period of time, we offer a complementary view compared to studies on best-practices. We investigate whether trust and social capital changed or remained stable over time in the offshoring project, and how these factors affected the offshoring relationship between the involved teams. In order to provide a detailed analysis, we investigated *articulation work* (Strauss 1988) conducted in the offshoring project.

The paper starts with a discussion of offshoring literature before it describes the research method applied in the ethnographic study. The description of the case is followed by a discussion under the perspective of articulation work. It turns out that trust and social capital indeed share some similarities, but are no guarantee for successful offshoring. Related findings are explained in the final chapter.

Offshore Cooperation in the Literature

Apart from challenges which are typical for any software development, GSE projects have to be conducted under particular organizational, cultural, spatial, temporal and legal conditions which can pose complex obstacles (Herbsleb et al. 2005; Ramesh et al. 2006). For example, temporal differences can lead to bottlenecks in regard to time for collaboration and coordination, while cultural differences can lead to mutual misunderstandings. As spatial distribution can harden or even constrain possibilities for control considerably, it often affects the necessary level of loyalty and trust among collaborators.

Hence, trust and social capital have been pointed out to be key factors for tackling challenges of distributed team cooperation (Hinds & McGrath 2006; Levina & Vaast 2008; Rottman & Lacity 2008). Trust has been characterized as a complex, multi-layered concept, which is—amongst others—related to expectations, experiences, and knowledge (e.g. is the trustee competent? Is his behavior predictable? Is he good-willing? Is he opportunistic?) (Imsland 2005). For our case, trust can be interpreted as a psychological state which allows for greater levels of self-organization, and for an abandoning of (available) control mechanisms (Zolin et al. 2004). In a similar fashion, social capital refers to network ties of “goodwill, mutual support, shared language, shared norms, social trust and a sense of mutual obligation that people can derive value from” (Huysman & Wulf 2004). As “social glue” holding together communities, social capital is expected to promote cooperative behavior in communities and organizations (Putnam 2000; Cohen & Prusak 2001).

According to this optimist view, organizations (or teams/communities within organizations) with high levels of social capital will have a higher motivation to cooperate (Huysman & de Wit 2004). However, with regard to offshoring, social capital can be difficult to be fostered (cf. Cramton & Hinds 2007; Hinds & McGrath 2006): as teams are distributed spatially, face-to-face contacts are usually reduced to few limited timeframes. At the same time, relying on ICT for cooperation implies a higher risk of misunderstandings (Olson & Olson 2000; Billings & Watts 2007), especially in cross-cultural teams. For the same reasons, conflicts can be very difficult to handle, if they occur (Hinds & Mortensen 2005).

As a consequence, recent studies have stressed the importance of initial perceptions of trustworthiness for long-term relationships of international teams (Lee et al. 2008). Inter-personal trust, once established, was found to remain more or less stable in the course of distributed projects, at least in case of cross-functional teams (Zolin et al. 2004). However, it is still not clear if the same rule applies to homogeneous fields with uni-functional dyads, such as in software-offshoring projects, where developers usually should be able to assess the development work done by the other team more easily as compared to cross-functional teams.

By revisiting the failure case story of a small German software company, we want to analyze if social capital shares the detected self-preserving effects of inter-personal trust relation, or if it may at least benefit from them. As trust can only be understood within a particular context (Zolin et al. 2004), it is necessary to take a situated perspective in analyzing the actual work practices, the distribution of labor as well as the related formal regulations and conflicts in offshore relations. In this context, the case of a failure story can be interesting if it allows for the differentiation between continuities and discontinuities.

However, continuities within formal regulations neither guarantee their factual continuity, nor does changing regulations guarantee factual discontinuity, as in reality established patterns may prevail under new labels, or formal regulations may fail. Hence, it is the practical organization of collaboration, not merely the mental models of its organizers, which has to be taken into account. As we were interested in long-term dynamics of offshoring software development, a methodological focus was needed that covers all above-mentioned aspects. This made us adopt the concept of *articulation work*, as we will point out in the following section.

Methodology

The concept of articulation work was introduced by the sociologist Anselm Strauss for the analysis of interdependent actions of cooperating actors (Strauss 1988). Articulation work, similar to coordination, is needed to regulate the division of labor: it centers on decision-making regarding who is supposed to do what, when, where, how, with which quality, etc.? To a certain degree, everybody involved in collaborative work has to reflect not only about his/her work, but also about its organization. In this regard, articulation work is also related to trust and social capital, as it entails issues of trading formal control versus flexible self-organization (Boden et al. 2008).

Generally, coordination is seen as the organization of collaborative work. However, not everything which is necessary for collaboration is explicitly discussed and regulated as coordination, and often the organization of work is more complex than perceived by many actors. For instance, collaboration may need meetings or discussions between developers, the management, and the customers, but it may also include the administration of a program for a certain task (setting up a related infrastructure), fixing a broken server, or implementing a communication infrastructure for collective work organization— aspects which are seldom interpreted as coordination.

In this regard, the concept of articulation work aims at including *all* necessary (meta-)work to make work work. Hence, it offers a more holistic understanding of cooperative work than concepts of coordination: while the latter usually govern the distribution of tasks and responsibilities, articulation work includes formal and

informal coordination mechanisms (Schmidt & Simone 1996) as well as related meta-work, which the actors themselves are sometimes not even aware of (Star & Strauss 1999).

As a consequence, it might not always be clear what should be regarded as meta-work or coordination when it comes to particular efforts. What one actor sees as necessary in this regard does not need to be the same as the perception of another. The same is true for scientific observers, who are influenced by their perception of the case. In this regard, coordination may be understood as the explicit model resulting from self-organization, and meta-work as the related practical conclusions, both of them being dependent on cognition and practical interpretation. In contrast, articulation work is the amount of all related contributions, strategies and conflicts; it is the distributed *agency* of collaboration, not its result.

Articulation work takes the individual perceptions about coordination neither as per se correct descriptions of the distribution of labor, nor as pure illusions; instead, they are understood as necessary points of departures for related analyses. Explicit (coordination) and practical premises (meta-work) of collaboration are regarded as important challenges for the individual positioning of actors within an anticipated field of opportunities. By contrasting conceptions of collaborators with each other and by analyzing empirical evidence on collaboration practices and their outcomes, articulation work studies attempt to take the interests of the collaborators seriously by discussing them retrospectively against the background of all accessible knowledge about the collaboration and its impacts.

This kind of analysis allows the address of the differences between the lived (factual) and the planned, explicit organization (cf. Argyris et al. 1985). The latter is generally more “logical” (at least at first glance) than the lived organization, which in contrast generally responds to situated particularities in a more complex way. This duality of formalized and informal organization has been discussed within the CSCW community (cf. Suchman 1987) for a long time and led to a much broader understanding of cooperative work in this community when compared to hierarchical models of coordination (Schmidt & Bannon 1992; Faergemann et al. 2005).

In order to study articulation work, we did not only have to look at efforts of coordination and meta-work, but also had to analyze them by contrasting the anticipated logic of the process with what we observed as the factual one. Revisiting the history of a cooperative project over several years with our theoretic stances is difficult, and requires careful examination. Unfortunately, our access to the company was limited to particular timeframes, and we had to reconstruct (and interpret) parts of the case study by relying to narrative interviews with the involved actors. However, we tried to overcome the limitations of our approach by a triangulation of several ethnographic research

methods, comprising semi-structured and narrative interviews, participant observation as well as artifact analysis.

In order to understand the logic of offshoring strategies, we started by collecting related conceptions in the literature. Furthermore, we conducted a semi-structured interview with the German manager of the company in 2006 that centered on his general perception of offshoring, as well as on his particular offshoring strategy.

For the investigation of articulation work practices, we drew on participant observations which were conducted by visiting the German SME several times during 2007 and early 2008. The first two observations lasted one week each and focused on local and distributed software development practices in individual and collaborative work situations and tasks. We were also allowed to analyze artifacts such as emails, chat protocols, internal work papers and whiteboard sketches, and we conducted many informal interviews with developers and the German project manager during our stay. The findings were documented by means of field notes and photos, which were taken during the research.

Our analysis of the collected data was based on Strauss' and Corbin's Grounded Theory (1998). After each step, the transcripts of the material, both field notes as well as interview data, were scrutinized. Data was coded during a process that consisted of several stages. At first, we composed categories based on the findings in the collected data. Then these categories were related to each other and evolved during the further research. These categories were analyzed under the presented articulation-work perspective. First, we attempted to differentiate between formal work organization (taken from the interviews) and the factual work practices we had observed. Then, we tried to identify converging and different perceptions of the offshoring project, as well as reconstructing related interests on the basis of a careful examination of our data.

As a further step, we refined the results of our analysis by conducting extensive narrative interviews with the German project manager during a third on-site visit, as well as a Skype interview with the Russian team manager.

The Case Study

The offshore software development project we researched was conducted by a German SME. The company offers a standard software solution for process modeling as well as services in the field of business process management. Being part of a holding, the management and sales of the company were handled by an office in Bonn with seven employees. The holding had several other offices, for example in Hamburg (data processing) and Düsseldorf (holding-management). Additionally, about 200 business consultants worked as freelancers in close cooperation with the company.

The software development we studied was carried out by an office in Berlin with seven employees. Apart from the development, the team in Berlin was accountable for the customer support as well as the management of the offshore cooperation with the Russian partner company. This cooperation had been started in 2002, when the German software office (at that time not yet integrated into the holding) had decided to found an offshore branch office in Saint Petersburg in order to reduce development costs.

The decision to locate the branch to Russia was based on a personal friendship of the German entrepreneur with a Russian developer who, according to the German manager, was trusted to be a competent and loyal team manager. This developer had been employed and ordered to hire three further developers in Saint Petersburg. The whole team was invited to Germany in order to become acquainted with the code base of the company. The team was able to take on the leading role of software development after a couple of months.

The German team manager described how the development of a formal model of work distribution marked the beginning of the cooperation. This model defined different roles and tasks for the teams. It included the role of the (German) project manager, the (Russian) team manager, the (Russian) software developers and the (German) testers. Thus, the German project manager wanted to oversee the development of the offshore team directly. In disciplinary or legal matters, the local team manager could be involved.

The German team concentrated on quality assurance, which involved the helpdesk for customers, the testing of the developed code and the strategic planning of upcoming versions. Thus, the definition of new features in terms of specifications and the description and classification of newly discovered bugs were under the responsibility of the German project manager and his team, while the offshore branch was responsible for the execution of the development. In the daily work, the results of the tests, descriptions of new features customers had asked for or bugs that were encountered by the helpdesk team and similar information would be communicated to the offshore branch for investigation. This usually involved personal visits of the German project manager to the offshore site shortly before new releases. During these visits, the German project manager helped handle the bugs (usually discovered in the last minute) and discussed the features of the following release with the Russians. The Russians in turn were to document their progresses in terms of monthly reports and review their code on a regular basis for quality assurance.

In the cooperation, members of both teams relied on several tools, which included a shared code repository (SourceSafe, situated in Germany) and IDEs/compiler (for C/C++, Java and Visual Basic), a bug tracking system as well as a product and development database based on Lotus Notes. For daily communication, a Lotus Notes plug-in called "Sametime" provided instant-messaging and screen-sharing functionality. Sametime allowed for the integration into the

Tool	Type	Used for
SourceSafe	CVS	Managing source-code
Product database	Lotus Notes database	Administrating specifications and releases; tracking progress of work
Development database	Lotus Notes database	Sharing templates for specifications, bug reports and formal work conventions
Sametime	Lotus Notes plug-in	Communicating via Instant Messages; sharing screens
SQA	Bug tracking system	Administrating and tracking bugs
Borland / Eclipse	IDE	Working on code

Table I. Tools provided for cooperation.

Lotus Notes environment and for encrypted communication and recording of screen-sharing sessions for later reviews (see table I).

Changes to the Division of Labor

According to the German team manager, the quick growth of the offshore team soon required certain adjustments of the formal division of labor. As he explained, it had become increasingly difficult to specify new features quickly enough to keep the growing offshore team busy—especially, when the number of Russian developers had exceeded the size of the German team. As the German project manager put it: “One day of development required one day of writing specifications” (Field notes, March 11, 2008).

As it became harder and harder for the German team to keep up with their work, the decision was made to change the formal division of labor. The Russian developers were now to write the specifications themselves, which were then in turn checked by the German team. According to a German team member, this decision was also based on the high competency of the Russian developers, who were trusted to have a deep understanding of the technical feasibility since they were in charge of the development. This change allowed the German team to reduce its work significantly and to enable the further growth of the offshore team which soon reached a size of up to 15 developers, as the German project manager reported.

However, delegating the requirement-engineering tasks to the offshore team led to significant problems, as the Russians lacked the necessary context knowledge: “The required information is very detailed: what does the user

interface look like, what conflicts may prevail with other features and what are special cases etc.” (Field notes, March 11, 2008). As the project manager pointed out:

“Knowledge concerning the practical usage and the technical background has to be combined in a creative way in order to find a solution. There is a difference between requirements specifications [considering the context-of-use] and design specifications, [being limited to the technical background]. The Russians tended to produce the latter” (Field notes, March 12, 2008).

According to the project manager, the problem was exacerbated by the Russian team’s poor English skills. While only rudimentary English skills would be needed for the coordination and control of already defined tasks, the definition of new features or the transfer of context information would be much more complex, thus sometimes exceeding the skills of the Russian colleagues. The German project manager explained: “The chats took much time and it was very difficult to transfer the related knowledge. It is easy to assign tasks or take over results, but it is hard to explain what needs to be done” (Field notes, March 11, 2008).

On the other hand, the Russians also reported problems concerning this way of cooperation:

“People [from the German team] (...) had no time to review them [the specifications], so the developers started to work without acceptance of specifications. (...) [So the] specifications did not follow the real implementation, or it took too much time for writing specifications” (Interview, May 28, 2008).

Attempts of Standardization

Faced with severe problems of communication and knowledge transfer, the company introduced a higher level of standardization to their documentation. Thus, standardized forms for documents, conventions for bug descriptions, source code comments and specific languages were developed. By providing examples and checklists, seen as help for the Russian developers with their tasks, the company expected to reduce the amount of communication and to ensure the quality of the produced documentation. The related documents were stored in the development database.

However, the complexity of writing specifications in combination with the missing background knowledge still made the tasks difficult and inconvenient for the Russian team, as the German project manager explained: “[The Russians] lacked the understanding of the program and the context of its use and the work is very unattractive, as it is very challenging and not well supported by tools” (Field notes, March 12, 2008). In addition, the German team reported increasing difficulties with the offshore developers, who started to ignore tasks that were

recognized by both teams as being unpleasant and annoying. This mainly included the writing of documentation and specifications as well as the tracking of the work with log-files.

“The Germans introduced forms to the product database in which the Russians should have entered their tasks with the expected beginning and end. They did this, but only at the beginning of the planning stage. As everything is very complex and unexpected dependencies occur, it is impossible to anticipate everything. Thus data needs to be updated regularly, but the Russians did not do so” (Field notes, April 12, 2008).

The following excerpt of a conversation illustrates this problem. The dialogue was taken from the chat-log of an online meeting between the German project manager and one of the Russian developers. The initiator of the online-meeting was the project manager who wanted an overview of the developer’s tasks. Using Sametime, the project manager was able to take control of the mouse and screen of the Russian developer and test the newly implemented features in this way. The inspection was accompanied by a chat discussion and took nearly three hours. The subjects of the discussion were the tasks (mainly feature specifications) in the product database, which were worked off feature by feature:

”Project manager I know we spoke MANY times about it... (...)
 it is impossible for me to follow progress if you don't write comments!
 so please don't let me repeat it again :-((...)
Developer I don't understand what should I write here, the implementation is
 fulfilled in 100 %
Project manager let me show you how i do it in my tasks.

On the shared screen, he [the German project manager] shows Dmitry some comments he has written. He opens one of his tasks, where he has already noted his progress like in a diary. The comments hint at problems he encountered, and at discussions with the developers. Then, he opens the product database and starts to comment on another task:

- started implementation
- bss [abbrev. name of another developer] send me new idea, so i stopped implementation. see info above.

Project manager you can decide the details in the comments.
 you should however add info that may be useful to you and to other
 people.
 this may help you keep notes on tasks (instead of using paper :-))
 or for example when you stop a task or need to restart it after some
 time.... you can use this to remember what you have to do.
 in any case whenever you update STATUS, PROGRESS or DATES....
 then you should add a comment regarding the reasons of the update.

Developer o.k.”
(Field notes, July 11, 2007).

Similar discussions concerned the conduction of internal code reviews of the offshore team as well as other examples of missing documentation.

Selling the Offshore Organization

According to the project manager, the problems with motivation were exacerbated when the decision was taken to sell the branch office to the Russian team manager in 2004. This decision was related to ongoing problems with the cooperation, as the Russian team manager reported:

“When we started [we were] four people (...). [In] 2004, all of these four developers left the company, because they were not satisfied with the situation. And from my side I wasn't able to do anything, to keep them (...). Because I had always to discuss any small question with Berlin” (Interview, May 29, 2008).

Furthermore, the decision was related to the challenges of handling the complex legal and organizational requirements of running an offshore branch. The communication with the local authorities turned out to be a serious and permanent challenge for the small German company, having no previous experience with Russian law. Hence, by changing the status of the Russian partner to that of an independent company, the German entrepreneur hoped to avoid many of the legal problems of managing an international company.

Thus, in 2004 the decision was made to continue the work by means of a contract between the SME and a now legally independent Russian company. As the Russian team manager explained, the Russians were quite happy with this change: “After we started to work as an independent company it got much easier for me to take decisions (...). And before, it took long discussions with Germany about why it was required” (Interview, May 29, 2008).

However, according to the German project manager, this change had dramatic consequences for the international cooperation. The Russian team manager, now being the proprietor instead of the employee, started to expand his company and look for new customers in order to reduce his dependence on the German SME.

In the interview, the German team manager described this strategy as expectable and even understandable. However, there were also unforeseen consequences, as it became much more difficult to continue the cooperation when the Russian team manager increasingly reduced his commitment to the cooperation. As finding new business partners became the main goal of the Russian partner, the German developers again were unable to control the Russian developers, who (from the perspective of the Germans) lacked discipline.

Even worse for them, according to the German project manager, the Russian team manager had been the most experienced and trusted team member abroad (especially since so many others had left the company), and his change of interest

led to severe problems, as the other Russian developers were unable to perform his duties with the same professional standard:

“[The German project manager] was unhappy that [the Russian team manager] was not available as a developer anymore from one day to another. As he was the manager instead of the developer now, the relationship had changed: instead of giving orders, everything was subject to negotiation” (Field notes, April 12, 2008).

In this regard, the dependency on the Russian developers made it difficult for the German team to enforce a reasonable accomplishment of tasks (especially of inconvenient ones) by the Russian team. “[The Russian developer] agreed to change his behavior, but he did not do it. And the Germans apparently were unable to convince him” (Field notes, April 12, 2008).

Salaries and Infrastructure

The problems with the offshore developers hit the company at a disadvantageous point of time. In 2006, the German company had been taken over by a holding. At the same time, according to the German project manager, the development costs had almost tripled compared to the situation in 2002. As both sides reported, the level of the salaries was an ongoing field of conflicts between the sides. As the Russian team manager explained:

“Finally they realized that they paid much more than they expected. (...) Salaries grew up too much in Saint Petersburg, and (...) I think, currently it makes not big sense to outsource from Germany to (...) Saint Petersburg. Because prices are comparable. (...) [And I told them] I was not ready to continue our contract on these terms” (Interview, May 28, 2008).

Because of the poor performance in combination with the rising development costs, the holding decided to reduce the size of the offshore team to eight—a decision, which further increased the frustration of the offshore team, as the German project manager reported.

According to him, the reduction in the number of employees belonging to the offshore team made it easier to coordinate the shared development, but the financial problems prevailed. He explained that this was due to the growing importance of Saint Petersburg as a software region. Western companies were in search for offshore developers, and the job market was growing rapidly. The lack of social security (sometimes seen as an argument for the attractiveness of a country) made income the only security for employees, and thus contributed to increased salaries. Policies of the German SME to keep salaries low were a constant field of conflict in the offshoring cooperation. At the same time, the small team size made the company especially vulnerable to fluctuation of team

members, while the low level of specialization required extensive training of new developers.

In this context, the German team reported that the Russians tried to use their influence on the development. Conflicts started about the distribution of (inconvenient) tasks and the technical infrastructure. For example, instead of using Sametime for their communication, the Russians started using Google Talk, and instead of using the company's Lotus Notes Database for shared documents, the Russians switched to Google Docs for their daily work.

Furthermore, the Russian team decided to stop using the shared bug-database SQA in favor of a self-developed database in 2007:

“The management of the company wants to get monthly reports concerning the ratio of feature development against the fixing of bugs. The tools [SQA and the product database] distinguish between both kinds, but it is not possible to [...] create an automated report, which the Russians find annoying. Therefore, they plan to administrate features and bugs in a shared database and have begun to develop their own, web based solution” (Field notes, July 10, 2007).

As a result, the teams had to track bugs in two parallel systems, because the German company was reluctant to change their established infrastructure. On the other hand, the German team manager did not want to antagonize the Russian team:

“Basically, [the German project manager] likes the idea [of a shared system], but the report feature is not necessary because they only need rough estimates for the taxes. But to avoid decreasing the motivation of the Russians they let them do as they like, as long it does not involve more work for the company” (Field notes, July 10, 2007).

Therefore, he did not intervene, but his acceptance was based on the condition that the Russians took on the necessary overhead work of maintaining two systems. In addition, the Russians planned to develop an import/export filter for the automatic synchronization of the two databases.

The other changes of development tools, i.e. using Google Talk instead of Sametime and Google Docs instead of Lotus Notes, were justified mainly with the available resources of the developers' computers. Since Sametime, according to the developers, needed much processor time and memory, it was annoying for the Russians to do their everyday work. Using the web-based Google Talk would be much more convenient for them. From the perspective of the German project manager, the decision had another reason. According to him, the Russians wanted to keep up to date with the tools they used. Thus Sametime and Lotus Notes would not be as trendy as the newer Google tools.

The Termination of the Cooperation

In 2007, the size of the offshore team was further reduced to four. Finally, in early 2008 the German holding decided to stop the cooperation completely, first by reducing the team size to two, and then by suddenly stopping the offshoring by the end of the month. The decision itself had neither been unexpected nor was it unwelcome by the German partners:

“All in all, everyone was unsatisfied with the state of affairs. The Russians, because the holding paid unpunctually, the developers in Berlin, because bad work was delivered, and the holding, because everything was considered as being too expensive, and the prices were increasing further” (Field notes, March 12, 2008).

Accordingly, both teams had considered the possibility of terminating the cooperation, and the German project manager had made up a plan together with the Russian team manager which was meant to arrange this termination to be as smooth and easy as possible for both teams. According to the project manager, this was not only due to his own team’s interests, but also due to the personal friendship with the Russian team manager. In this regard, both teams said they would have liked to continue the cooperation under different circumstances, and they blamed the holding management as being the one responsible for the failing of the project.

Hence, in the end, only the abruptness of the decision caught both teams by surprise. As the Russian team manager explained:

“In the middle of December, [the German holding] said, o.k., please keep these four developers until end of May (...). So we will have five months to move the development from Saint Petersburg to Germany. (...) But [then] they said that they had changed their decision and needed only two people until the end of February. This was unexpected (...) and I had to pay salaries for them and even (...) fire one developer“ (Interview, May 29, 2008).

Analysis of Articulation Work and Social Capital

While the last chapter recapitulated the course of events from the perspective of the practitioners, we will now revisit the offshoring story from an articulation work and social capital perspective.

As we were told, the initial phase of the cooperation was supported by a high level of trust between the teams, which was based on the friendship between the German entrepreneur and the Russian team manager. Furthermore, the visit of the whole Russian team to Germany had helped to form social ties between the developers, too. However, despite this high level of social capital, the German team wanted to stay in control of the development as much as possible, as software development was still deemed as the core competency of the company.

The Russians, on the other hand, accepted this distribution of tasks, as it allowed them to concentrate on the technical side of the development only.

However, in order to do so they were dependent on exhaustive specifications of features which the Germans found increasingly difficult to afford. As the initial distribution of labor turned out to be problematic, the German company had to learn that writing complete specifications (even for the standard software product) can be as time-consuming as the development itself (or even more). Instead of being self-explanatory and efficient, the disjunction of requirements-engineering and coding led to severe coordination problems which were caused by the necessary knowledge transfer and articulation work between the teams.

As the workload of the German team increased, the decision was taken to change the distribution of work while ensuring that control remained with the German team. The Russians accepted this change unwillingly. Despite their good technical knowledge, the Russians had difficulties with the task of writing specifications. As they lacked the necessary context knowledge, the effort of writing adequate documentation was very high, even more as they could not draw upon shared business experience with the customer. As a result, the Russians felt overstrained, and the amount of necessary requests, clarifications, and corrections increased—classical aspects of articulation work. At the same time, the dependency on ICT for articulation work created bottlenecks, which were further aggravated due to language issues between the teams. While the trust between the teams was still high, the German team attempted to improve the documentation by introducing standardized forms for specifications and bug descriptions. However, this attempt to support the Russian team in writing specifications did not work, as the necessary knowledge exchange was still insufficient. In contrast: the efforts to formalize the development turned out to be only new forms for informal articulation work, and for related uncertainties in the development process.

Despite the related increase of informal communication—which has been found to support knowledge exchange and even conflict resolution in distributed teams (Hinds & Mortensen 2005)—the company could not benefit from the change, because the Russians lacked the necessary context information which was paramount to successfully accomplish the task of writing proper specifications. After all, communication needs to support the underlying work structures of a team (Hinds & McGrath 2006). Instead of supporting the necessary communication work, the management of the holding—bound by the necessity to coordinate two organizations—reacted by intensifying control and formalization. This in turn was seen as an escalation and systematization of attempts to blame the Russians for the prevailing problems. The social capital which had formed the basis of the commitment of the Russian colleagues started to become eroded—despite the initial high level of trust between the sites, which rested on the personal relationship between the German project manager and the Russian team manager.

The Russians, unable to meet the expectations of the German team, began to neglect certain tasks which were regarded as being annoying and unnecessary, like tracking the progress of their work. The German company had to realize that it was dependent on the commitment of the Russian team and that formal methods of control cannot guarantee personal obligations—or even damage them (cf. Imsland 2003). Even worse, the German team was unable to solve this problem. In this phase of the cooperation, the still high level of social capital apparently hindered an open argument between the teams. The German management avoided blaming the Russians outrightly for not fulfilling their tasks, while the Russians avoided arguments with the German side by simply ignoring inconvenient tasks. In this regard, social capital apparently became a trap: the German manager understood the anger of the Russian team, but regarded the current division of labor as necessary. The Russians, on the other hand, accepted the decisions of the Germans, but felt unable to work under these conditions.

As more and more of the Russian developers left the company, the decision was taken to sell the offshore organization to the Russian team manager. While this decision was approved by both sides, it became the origin of further emerging conflicts, as the cooperation with the now legally independent Russian enterprise made it impossible for the Germans to use hierarchy in order to maintain their idea to substitute informal demands of articulation work by means of intensified formalization. The loss of competent developers was a significant drawback for the company, not only in terms of knowledge, but also in terms of social capital. While the initial cooperation had rested on the personal ties which were formed during the extended personal visit of the Russian team to Germany, the newly hired developers could not benefit from such relations. Instead, the social ties between the teams were mainly focused on the Russian team manager, who shifted his focus to acquire new customers instead of concentrating on the existing cooperation. As a result, the problems with the motivation of the Russian developers aggravated, as social capital as a means for motivating cooperation between the German team and the new Russian developers was weak.

The German team—discontent with the development of the cooperation with Russia—felt trapped: since none of the teams was able to work efficiently without the other team, but every team had the possibility to jam shared projects (by ignoring or by misinterpreting cooperation demands), successful cooperation became unlikely. Collaborative demands on articulation work—considered to be substitutable by formalization and control by the holding management—emerged again on each level of conflict resolution and turned out additional strategic options for constraining the cooperation afterwards.

In this regard, the company also suffered from hard-to-anticipate indirect effects, as, for example, the rising salaries in the region of Saint Petersburg, which were partly connected to the boom of investments in the area and contributed to cost the inefficiency of the cooperation and its termination. Moreover, the change

of infrastructure on behalf of the Russians was hard to foresee. The German team accepted these changes within certain limits because they feared to further discourage the offshore team members.

Apparently, the management of the holding overestimated the possibilities of formal control, and neglected conflict dynamics and social capital issues. As a result, conflicts manifested when coordination necessities emerged on the basis of inter-dependencies in the work constellation which could not be settled by controlling and formalizing the software development. When people tried to solve the problems by means of formalizing articulation work, the situation did not improve, but deteriorated—and was further aggravated by the structural circumstances like rising costs, decreasing social capital and the organizational consequences of the divestment of the branch office.

Conclusion

Our case study illustrates the endeavors of a small German enterprise to keep its offshoring project running. Looking at the related failure story from a long-term perspective, complex and inter-related conflicts within the field of articulation work become visible. The German management, for instance, tried to take advantage of the relatively low wage levels of the Russian partners whom they sought to control by determining their work (by means of the division of labor and tool usages).

The Russian partners turned this claim around by showing that, if control was that important, there was a lack of it in the whole collaboration. How could the Russians work well if the requirement delivered to them were not controlled (for instance, if they complied with international standards)? The counter-reaction of the German management was, again, a turn-around: if the requirements were that difficult to handle, the Russians should write them themselves.

This shows that decisions were taken to shift responsibilities between the offshoring partners. Therefore, articulation work obviously was not only a contingent dimension of decision making, but (in the given case) even attributed to a history of its own in regard of work regulations. The related ping-pong effect of control-based arguments shows that both partners shared related convictions or, at least, did not want to question them. Hence, the changes to the work arrangements were partly the result of continued shared convictions about the necessities of control and formalization. But those were not the only continuities.

While it became apparent that any new regulation led to new areas of conflict, it has to be noted that this did not diminish the mutual appreciation among the actors. Their mutual trust remained through all these conflicts. But what about social capital? It was defined before as “network ties of goodwill, mutual support, shared language, shared norms, social trust, and a sense of mutual obligation that people can derive value from” (Huysman & Wulf 2004). Have these ties declined

throughout the diverse conflicts? The astonishing fact is that partners from both sides still would have liked to collaborate even after the termination of the offshoring project which was seen as a salvation on both sides.

Obviously the actors differentiated between the personality of their partners and the offshoring situation as a whole (cf. Imsland 2005). This implies that the management partly understood the strategies of the Russians to take advantage of higher salaries, or at least did not see it as personally insulting. In contrast, the Russians obviously understood the role of the holding as a limiting factor for the German project manager. Insofar, the trust—with regard to mutual goodwill—among the actors prevailed even through disappointing collaborative experiences regarding opportunistic and sometimes unpredictable behavior. The same was true for social capital in the mentioned sense as an accumulative value.

However, the social capital, which contributed to motivation at the start of the offshoring project, also turned out to be a hindrance at its end, as the assumed knowledge about the personalities of the partners made it easier to detect structural limitations of the situation, and as it apparently hindered an open argument about the prevailing conflicts. This means that social capital can really be an asset in the sense that collaboration would not be possible without it. Nevertheless, it can become dysfunctional, a mis-investment in terms of the capital metaphor. Social capital is not only about cognition, inter-action, and shared perceptions: it also relates to fallible investment of efforts.

This fallible characteristic of social capital is not covered in the necessary detail by Putnam's tradition of social capital as "goodwill, mutual support, shared language, shared norms, social trust, and a sense of mutual obligation". Hence, it seems to be fruitful to expand the given understanding of social capital by referring to Bourdieu's (1986) notion of social capital as a means to reconstruct the risky decisions of individuals when attempting to establish profitable value chains, which can explain why social capital apparently can change from an asset to a hindrance.

In relation to offshoring, it was found that there seems to be something like a tendency of trust to prevail (Zolin et al. 2004). We came to similar results for social capital, but our results also question the concept of social capital as a merely positive factor for global software engineering. Like for trust, it seems we need a much more differentiated understanding of social capital in the context of GSE. Without social capital, GSE as a complex form of distributed collaboration will hardly be possible. On the other hand, formalization and social capital are no guarantee for successful performance. As we have seen, impacts of the international environment and contingencies of articulation work make it very likely in GSE that a given arrangement changes quickly. Therefore, it seems to be a major challenge for GSE to develop forms of making articulation work reflexive and operative, for example, through globally distributed organizational learning.

Acknowledgements

This paper was supported by the German Research Foundation (DFG) within the project “ARTOS – Articulation work in offshoring projects of small and medium-sized enterprises of the software branch”.

References

- Argyris, C., Putnam, R., and Smith, D. M. (1985). *Action science*. Jossey-Bass, San Francisco.
- Billings, M. & Watts, L. (2007). ‘A Safe Space to Vent: Conciliation and Conflict in Distributed Teams’. *European Conference on Computer Supported Cooperative Work*, Limerick, Ireland.
- Boden, A., Nett, B., and Wulf, V. (2008). ‘Articulation work in small-scale offshore software development projects’ *International workshop on Cooperative and human aspects of software engineering*, Leipzig, pp. 21-24.
- Bourdieu, P. (1986). ‘The form of capital’, in J.G. Richardson (ed.), *Handbook of theory and research for the sociology of education*, Greenwood Press, New York.
- Cataldo, M.; Wagstrom, P. A.; Herbsleb, J. D. and Carley, K. M. (2006). ‘Identification of Coordination Requirements: Implications for the Design of Collaboration and Awareness Tools’, *International Conference on Computer Supported Cooperative Work*, pp. 353-362.
- Cohen, D. and L. Prusak (2001). *In good company: how social capital makes organizations work*, Harvard Business School Press, Boston.
- Cramton, C. D. and Hinds, P. J. (2007). ‘Intercultural Interaction in Distributed Teams: Salience of and Adaptations to Cultural Differences’, in G. Salomon (ed.), *Proceedings of the Academy of Management Annual Meeting, Best Papers*, Philadelphia.
- Faergemann, L., Schilder-Knudson, T. and Carstensen, P. H. (2005). ‘The Duality of Articulation Work in Large Heterogeneous Settings – a Study in Health Care’. *European Conference on Computer-Supported-Cooperative-Work*, pp. 163-183.
- Fisher, J., Hirschheim, R., and Jacobs, R. (2008). ‘Understanding the outsourcing learning curve: A longitudinal analysis of a large Australian company’. *Information Systems Frontiers*, 10(2), pp. 165-178.
- Gutwin, C.; Penner, R. and Schneider, K. (2004) ‘Group Awareness in Distributed Software Development’, *International Conference on Computer Supported Cooperative Work*, pp. 72-81.
- Herbsleb, J. D.; Finholt, T. A. and Grinter, R. E. (2001). ‘An Empirical Study of Global Software Development: Distance and Speed’, *International Conference on Software Engineering*, pp. 81-90.
- Herbsleb, J. D.; Mockus, A.; Finholt, T. A. and Grinter, R. E. (2000). ‘Distance, dependencies, and delay in a global collaboratio’, *International Conference on Computer Supported Cooperative Work*, pp. 319-328.
- Herbsleb, J. D.; Paulish, D. J. and Bass, M. (2005). ‘Global Software Development at Siemens: Experience from Nine Projects’, *International Conference on Software Engineering*, pp. 524-533.
- Hinds, P. & McGrath, C. (2006). ‘Structures that work: Social structure, work structure, and performance in geographically distributed teams’. *International Conference on Computer Supported Cooperative Work (CSCW)*, Banff, Canada.

- Hinds, P. & Mortensen, M. (2005). 'Understanding conflict in geographically distributed teams: An empirical investigation'. *Organization Science*, 16, 290-307.
- Huysman, M. and D. de Wit (2004). 'Practise of Managing Knowledge Sharing: Towards a Second Wave of Knowledge Management', *Knowledge and Process Management*, vol. 11, No 2, pp. 81-92.
- Huysman, M. and V. Wulf (2004). *Social capital and information technology*, MIT Press, Cambridge, Mass.
- Imsland, V. (2003): *The Role of Trust in Global Software Outsourcing Relationships*. Ph.D. thesis, Univ. of Oslo.
- Jarvenpaa, S. L. and Leidner, D. E. (1995). 'Communication and Trust in Global Virtual Teams', *Journal on Computer-Mediated Communication*, vol. 3 (4).
- King, W. R., and Torkzadeth, G. (2008). 'Information Systems Offshoring: Research Status and Issues', *MIS Quarterly*, vol. 32/2.
- Lee, J., Huynh, M., & Hirschheim, R. (2008). 'An integrative model of trust on IT outsourcing: Examining a bilateral perspective', *Information Systems Frontiers*, vol. 10(2), pp. 145-163.
- Levina, N., & Vaast, E. (2008). 'Innovating or Doing as Told? Status Differences and Overlapping Boundaries in Offshore Collaboration', *MIS Quarterly*, vol. 32 (2).
- Olson, G. M. and Olson, J. S. (2000). 'Distance Matters'. *Human-Computer Interaction*, vol. 15, pp. 139-178.
- Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of the American Community*. Simon & Schuster, New York.
- Ramesh, B.; Cao, L.; Mohan, K. and Xu, P. (2006). 'Can distributed software development be agile?', *Communications of the ACM*, vol. 49, pp. 41-46.
- Rottman, J., and Lacity, M. (2008). 'A US Client's learning from outsourcing IT work offshore', *Information Systems Frontiers*, vol. 10(2), pp. 259-275.
- Schmidt, K. and Bannon, L.(1992). 'Taking CSCW Seriously: Supporting Articulation Work', *Computer Supported Cooperative Work*, vol. 1 (1992), pp. 7-40.
- Schmidt, K. and Simone, C. (1996). 'Coordinaton Mechanisms: Towards a Conceptual Foundation of CSCW Systems Design', *Computer Supported Cooperative Work*, vol 5, pp. 155-200.
- Sole, D. and Edmondson, A. J. (2000). 'Knowledge Sharing in Virtual Teams', *Harvard Business School Working Paper*
- Star, S. L. and Strauss, A. L. (1999). 'Layers of Silence, Areas of Voice: The Ecology of Visible and Invisible Work', *Computer Supported Cooperative Work*, vol. 8, pp. 9-30.
- Strauss, A. L. (1988). 'The Articulation of Project Work: An Organizational Process', *The Sociological Quarterly*, vol. 29 (2), pp. 163-178.
- Strauss, A. L. and Corbin, J. M. (1998). *Basics of qualitative research: techniques and procedures for developing grounded theory*. Sage Publications, Newbury Park.
- Suchman, L. A. (1987). *Plans and situated actions: the problem of human-machine communication*. University Press, Cambridge.
- Zolin, R., Hinds, P., Fruchter, R. & Levitt, R. (2004). 'Interpersonal trust in cross-functional, geographically distributed work: A longitudinal study', *Information & Organizations*, vol. 14, pp. 1-26.

Return On Contribution (ROC): A Metric for Enterprise Social Software

Michael J. Muller*, Jill Freyne**, Casey Dugan*, David R. Millen*,
and Jennifer Thom-Santelli***

*IBM Research, Cambridge, MA 02142 USA, **CLARITY, University College
Dublin, Dublin 4 Ireland, ***Cornell University, Ithaca, NY 14850 USA
*{michael_muller, cadugan, david_r_millen}@us.ibm.com, jill.freyne@csiro.au,
jt17@cornell.edu*

Abstract. The value of enterprise social media applications, components, and users is difficult to quantify in formal economic terms such as Return On Investment. In this work we propose a different approach, based on human service to other humans. We describe a family of metrics, Return On Contribution (ROC), to assist in managing social software systems. ROC focuses on human collaboration, namely the creation and consumption of information and knowledge among employees. We show how ROC can be used to track the performance of several types of social media applications, and how ROC can help to understand the usage patterns of items within those applications, and the performance of employees who use those applications. Design implications include the importance of “lurkers” in organizational knowledge exchange, and specific types of measurements that may be of value to employees, managers, and system administrators.

Introduction

This short paper proposes a new measurement concept and initial quantification to measure the business benefits of social software applications. Rather than focus on financial advantages, which are typically very difficult to measure for social applications, we propose to emphasize the human benefits of systems that link workers in diffuse networks of mutual aid.

Many approaches to the evaluation of commercial systems rely on the concept of Return On Investment, or ROI (e.g., Webb, 2008). ROI is sometimes easy to measure, when for example the application fills a crucial, measurable business need. ROI is more difficult to measure for applications or tools that operate in a more diffuse or supporting function (Howlett, 2007; Webb, 2008).

For example, it is straightforward to measure the cost to provide a telephone on each employee's desk. For a few select jobs (e.g., call centers), it is possible to quantify the business value of the telephone. But can an organization measure the value of that telephone for the rest of its employees (e.g., Howlett, 2007)? Is it appropriate to monitor the *usage* (calls made and received)? Is it appropriate to count *displacements* against other media – e.g., the number of physical letters not mailed? The contribution of the telephone itself is more difficult to measure than the many business functions of the employee who uses that telephone.

In this short paper, we propose a different approach. We define a set of measurements based on the concept of Return On Contribution, or ROC. Like ROI, ROC is a measure of benefit divided by cost. Unlike ROI, ROC focuses on human workers, and can be applied in situations without direct monetary metrics. Also unlike ROI, ROC emphasizes both the production of knowledge and the consumption of knowledge (as defined locally by human actors). In this way, ROC is part of the resurgence in interest in the subtle contributions of lurkers to their organizations (Nonnecke and Preece, 2001; Takahashi et al., 2003).

For the remainder of this note, we present a first definition of ROC, and show how that metric can be used to describe the overall *human* benefit of two enterprise social software applications. We then show how this concept can be focused on particular components and particular actors within such systems. We close with recommendations for design.

Return On Contribution

Return On Contribution is a ratio of benefit divided by cost. The “units” of the metric are persons. Within the framework of rational choice theory (see Pirolli, 2007, for a recent summary), we assume that employees make appropriate and strategic use of available collaborative resources. We therefore count each access to a resource a measure of the subjectively-defined value of the resource by the person who accessed it, and thus as an indirect measure of benefit to that person.

The core definition of ROC is the ratio of the number of people who benefit in this way from a resource (i.e., through rational consumption of that resource), divided by the number of people who create or contribute to that resource. For a social-media application, we can operationalize this definition by characterizing users as *originators* of the resources in the system, or as *consumers* of those resources. The primary focus of this project is to provide ROC as a metric of social value for social software systems in which users take actions to contribute

Measure	Social- Bookmarking	Social- Networking
Consumers	10896	21453
Originators	4213	8397
ROC_C = Consumers/Originators	2.59	2.55
Originators-&-Consumers	3654	7987
Lurkers (Consumers-only)	6683	13466
Originators-only	559	410
ROC_L = Lurkers/Originators	1.59	1.60

Table 1. Calculating Return On Effort. ROE_C is the overall ROE for all Consumers of information, in which some Consumers may also act as Originators. ROE_L is a revised figure based primarily on Lurkers (users who consume but never originate).

content, and other users receive value by receiving that content. Examples of such systems include social bookmarking sites, wikis, blogs, and file-sharing services.

Our analysis is different from the “authors vs. readers” approach of Noll and Meinel (2007). Their study compared the “authors” formal metadata (in HTML structures and internet rating systems) vs. the bookmarks created by “readers.” In our language, all of their users were *originators* (“authors” originate documents, while “readers” originate bookmarks), and there were no data about *consumers* who created neither documents nor bookmarks.

ROC on Entire Applications

Although ROC has a broad scope, here we applied the concept to two enterprise social media applications for which we had usage data (Millen et al., 2006; DiMicco et al., 2008) (Table 1).

- During July 2005 - April 2007, a social-bookmarking application contained contributions by 4213 bookmark-originators, and was directly used (consumed) by 10896 bookmark-readers. For this application $ROC = 10896/4213 = 2.59$ consumers of the work of each originator.
- During June 2007 - January 2008, a social networking application contained contributions by 8397 item-originators, and those contributions were viewed (consumed) by 21453 viewers. For this application $ROC = 21453/8397 = 2.55$ consumers of the work of each originator.

Table 1 shows examples of several variants on the ROC concept which are further refined by examining how many users act in the role of both originator and consumer. The top of the table shows the calculation of ROC_C (measured in terms of *all* Consumers), as described above. The bottom of the table shows the calculation for ROC_L (measured in terms of Lurkers *only*) – i.e., an ROC based on “pure” consumers who never explicitly contribute. These summary indices can be




Figure 1. ROC over time for two social-media applications.

used to show the spreading benefit of social software, from a core group of originators to a much larger group of consumers. In the remainder of this paper we will focus on the ROC_C measures.

ROC Over Time

ROC can also be calculated on a temporal basis for an application, to support the examination of growth and change over time. An administrator might monitor the organizational value of a social software application through ROC. Changes in ROC might indicate barriers to usage, and could be used to sense opportunities to intervene so as to enable or facilitate greater participation and system adoption.

Figure 1 shows the growth in ROC for both of the social media applications mentioned above, during the respective study period for each application. The social-bookmarking application appears to have begun robustly, with an ROC over 1.0 during the first month. While the monthly figures are somewhat variable, the generalized ROC_C never dips below 2.0 consumers/originator, and even the more refined ROC_L never goes below 1.4. This is to be evidence of a relatively stable pattern of use. In the terms of rational choice theory, the social-bookmarking application appears to benefit both originators and consumers.

The social-networking application shows a different pattern. Since its initial deployment, it has experienced viral growth (DiMicco et al., 2008) as shown by the nearly monotonic increase in ROC measures over the first seven months of deployment. This application does not yet appear to have achieved a “steady state,” so the administrator may look forward to even stronger patterns of usage.

ROC on Application Components and Persons

While ROC can provide an overall picture of benefit, administrators may want to uncover specific information about components of a social system that are driving the benefit. Are all media types and specific objects used with equal benefit? Are




Figure 2. ROC_C for different components of the social-networking application. Events were introduced in September 2007

all contributors comparable? Variants of ROC can be targeted for detailed analysis of components and persons.

The social-networking application contains three major types of media whose usage can be measured on a per-item basis: Photos, Lists, and Events (DiMicco et al., 2008). Figure 2 shows the monthly ROC_C metric for photos, lists, and events.¹ There were small, suggestive upward trends in ROC for Photos and Lists from June-August, but the major increases in ROC occurred for the Event objects. An administrator – or the leader of an online community – might want to study the Event genre to determine which of its attributes led to so much user uptake.

The domain of social tagging offers additional opportunities to use ROC in more fine-grained analyses. As described in Ames and Naauman (2007), some content-originators use tagging to reach large audiences. In the enterprise domain, employees reach large groups of colleagues through the strategic use of social-tagging in roles such as “evangelist” or “publisher” (Thom-Santelli et al., 2008) or “information curator” (Muller et al., 2009).

We can conduct analyses of the ROC of specific tags, by counting the number of people who include each such tag while they are creating bookmarks, and by counting the number of people who search on each such tag to find bookmarked content. These analyses allow us to find tags with relatively high ROC. For example, an “evangelist” was promoting awareness of social-media using the tag “web2.0”, and that tag had a per-tag ROC of 1.95 consumers/originator. Similarly, a periodic internal podcast “publication” was bookmarked by its authors with the tag “Tag-City”, and that tag had a per-tag ROC of 7.41.

We can also find specific creators whose tags are searched by a large number of their colleagues. In the previous paragraph, the “web2.0” tag was searched by many information-consumers, so the “evangelist” user who communicated

¹ The curves for ROC_L were very similar. We omit those results to conserve space.

Service	Consumers	Originators	ROC _C
Wiki server	238838	36377	6.57
Discussion server	150000	23000	6.52
People-tagging application	20973	3102	6.76
File-sharing service	68762	11276	6.19

Table 2. ROCs for four enterprise social software services.

through that tag had a personal ROC of 1245.0. Similarly, the “Tag-City” tag was searched by a large number of information-consumers, so each of the two “publisher” authors of the podcasts had a personal ROC of 63.0. These ROC values are strong evidence of the contributions of the originators to their consumers and their organizations.

ROCs for Other Enterprise Social Software Services

We obtained summary data from four additional enterprise social software applications that were beyond the scope of our detailed study. The provisional ROCs for these services were all greater than 6.0 consumers/originator (Table 2).

Summary

We have shown that ROC can be used to assess an entire system, and to track the usage of that system over time. We have shown how ROC can be applied to specific types of objects in a system, and we believe that ROC can also be used to compare the organizational significance of genres of objects (e.g., photos, lists, and events in Figure 2). Finally, we have shown briefly that ROC can be used to compare specific points of articulation (e.g., tags) within a social media application, and can also show the service of particular employees to their colleagues.

Looking toward the future, we envision more ways to use ROC. This paper examined the ROC of applications in which users make *explicit* contributions of content or ratings. ROC can also be an effective measure in systems that are purely lurker-driven. For example in Collaborative Web Search (Freyne and Smyth, 2006), users’ search activities are interpreted by the application as relevance judgments, and are displayed to assist subsequent users with similar searches. Because all users are, by definition, both explicit consumers (they search) and implicit contributors (their searches produce useful data), the ROC of such systems would be always be 1.0. By contrast, the systems that we studied in this paper have ROCs in the range of 2.0-3.0, and the systems summarized in Table 2 have ROCs in the range of 6.0-7.0. With more experience, we may be able to describe “characteristic ROCs” for different genres of social media.

Implications for Design

The ROC metrics depend crucially on measures of information-consumption, as well as information-creation or origination. CSCW systems have tended to focus on the creators of information, and to leave the consumers unmeasured – or to dismiss consumers as “lurkers” or “free-riders.” Indeed, lurkers have often been considered to be a problem because they consume but do not contribute – an issue that has been discussed in the language of the “tragedy of the commons” (for review, see Curien et al., 2006; Kollock and Smith, 1996). By contrast, certain web metrics have begun to highlight the importance of consumers’ behaviors for website maintenance (Saleem, 2008) or marketing (Fox, 2007; Webb, 2008).

With our ROC metric, we join Nonnecke and Preece (2001) and Takahashi et al. (2003) in the re-evaluation of the role of lurkers, especially in an organizational context. Nonnecke and Preece reported that some lurkers lurk for altruistic, pro-social reasons. Takahashi et al. showed that some lurkers use the information they have found to make contributions in ways other than the creation of entities in software applications. Enterprises often designate employees whose job involves the *origination* of knowledge and information, and other employees whose job involves the responsible *consumption* of that knowledge and information. These employees who are, in effect, “paid to lurk” perform valuable work for their organization and, often, for their clients and customers. The ROC metric focuses on measuring how these lurkers consume that information, and thus helps to highlight the importance of lurkers in organizational performance.

These observations lead us away from “tragedy,” and toward a “celebration of the commons.” Specifically:

- Social media applications should record and analyze the activities of information-consumers, not only to improve performance and to extend their marketing, but also to understand what information and knowledge is proving to be valuable to employees, and to tune the resources and their distribution to improve the sometimes *mission-critical lurking* of these employees. Summary statistics across groups of lurker-workers can help organizations to highlight the most important resources for those workers.
- Social media applications should allow administrators to track the ROC of the application as-a-whole.
- Social media applications should allow information originators to examine the ROC of the items that they originate, so as to evaluate and manage their effectiveness in reaching their intended audience.
- Managers of information-originators may wish to examine the ROC of the resources produced by their employees. More controversially, managers may also wish to examine the per-employee ROC. This idea is common in journalism, publishing, and information services. It remains to be seen whether this concept can also become part of the organizational

recognition of the contributions of knowledge work and knowledge-workers.

References

- Ames, M. & Naaman, M. (2007). 'Why we tag: motivations for annotation in mobile and online media.' *Proc. CHI 2007*, ACM Press, San Jose, CA, USA, April 2007, 971-980.
- Curien, N., Fauchart, E., Laffond, G. & Moreau, F. (2006). 'Online consumer communities: Escaping the tragedy of the digital commons.' In E. Brousseau (Ed.), *Internet and Digital Economics*. Cambridge: Cambridge University Press, 2006.
- DiMicco, J., Millen, D.R., Geyer, W.G., Dugan, C. Brownholtz, B., & Muller, M.J. (2008), 'Motivations for social networking at work.' *Proc CSCW 2008*, ACM Press, Banff, AL, Canada, November 2008, 711-720.
- Freyne, J., & Smyth, B. (2006), 'Cooperating search communities.' *Proc Adapt. Hypermedia & Adapt. Web-Based Systems*, Springer, 101-110.
- Fox, M., (2007) 'Social media ROI.' <http://socialmediagroup.ca/2007/11/13/social-media-roi/>, 13 November 2007.
- Howlett, D. (2007). "'ROI is so Business 1.0: not.'" <http://blogs.zdnet.com/Howlett/?p=183>, 1 Oct 2007.
- Kollock, P., & Smith, M. (1996). 'Managing the virtual commons: Cooperation and conflict in computer communities.' In S. Herring (ed.), *Computer-Mediated Communications: Linguistic, Social, & Cross-Cultural Perspectives*. Amsterdam: Benjamins (1996), 109-128.
- Millen, D.R., Feinberg, J., & Kerr, B. (2006). 'Dogear: Social bookmarking in the enterprise.' *Proc CHI 2006*, ACM Press, Montréal, QU, Canada, April 2006, 111-120.
- Muller, M.J., Millen, D.R., & Feinberg, J. (2009). 'Information curators in an enterprise file-sharing service.' *Proc ECSCW 2009*, Springer-Verlag, Vienna, Austria, Sep. 2009, in press.
- Noll, M.G., & Meinel, C. (2007). 'Authors vs. readers – A comparative study of document metadata and content in the WWW.' *Proc. DocEng'07*, ACM Press, Winnipeg, MA, Canada, August 2007, 177-186.
- Nonnecke, B., & Preece, J. (2001). 'Lurker demographics: Counting the silent.' *Proc CHI 2000*, ACM Press, Den Hague, Netherlands, April 2000, 73-80.
- Pirolli, P., *Information Foraging Theory: Adaptive Interaction with Information*. New York: Oxford, 2007.
- Saleem, M. (2008). 'Social media marketing ROI – Metrics and analytics.' <http://searchengineland.com/080826-123600.php>, 26 August 2008.
- Takahashi, M., Fujimoto, M., & Yamasaki, N. (2003). 'The active lurker: Influence of an in-house online community on its outside environment.' *Proc GROUP 2003*, ACM Press, Sanibel Island, FL, USA, November 2003, 1-10.
- Thom-Santelli, J., Muller, M., & Millen, D. (2008). 'Social tagging roles: Publishers, evangelists, leaders.' *Proc CHI 2008*, ACM Press, Florence, IT, April 2008, 1041-1044.
- Webb, D. (2008). 'Measuring intangibles revisited – Social media metrics and ROI.' <http://missiondrivenmarketing.com/2008/02/21/measuring-intangibles-revisited-social-media-metrics-roi/>, 21 February 2008.

Collaborative Practices that Support Creativity in Design

Dhaval Vyas, Dirk Heylen, Anton Nijholt

Human Media Interaction, University of Twente, the Netherlands

d.m.vyas@cs.utwente.nl

Gerrit van der Veer

School of Computer Science, Open University, the Netherlands

Abstract. Design is a ubiquitous, collaborative and highly material activity. Because of the embodied nature of the design profession, designers apply certain collaborative practices to enhance creativity in their everyday work. Within the domain of industrial design, we studied two educational design departments over a period of eight months. Using examples from our fieldwork, we develop our results around three broad themes related to collaborative practices that support the creativity of design professionals: 1) externalization, 2) use of physical space, and 3) use of bodies. We believe that these themes of collaborative practices could provide new insights into designing technologies for supporting a varied set of design activities. We describe two conceptual collaborative systems derived from the results of our study.

Introduction

A typical design studio (professional or academic) has a high material character – in the sense that it is full of material objects and design artefacts; office walls and other working surfaces full of post-it notes, sketches and magazine clips for sharing ideas and inspiration; physical models and prototypes lying on the desks and so on. The physical surroundings of a design studio and the persistence with which different material artefacts are arranged and represented are important to the design activity and serve as organizational memory (Ackerman and Halverson, 1995) and distributed cognition (Hutchins, 1995) for design teams.

This ecological richness of design studios stimulates creativity in a manner that is useful and relevant to the ongoing design tasks. Additionally, designers do not work in a stereotypical or mechanical fashion when designing interactive products. Designers tend to be innovative, creative and often playful in order to collaborate and successfully meet the demands of building new products and services. Methods frequently used by designers such as role playing (Boess, 2008), body storming, design choreography (Klooster and Overbeeke, 2005) and so on are not limited to problem solving but also include understanding interactional and experiential qualities in designing interactive products.

The role of collaboration between co-designers is critical to a design studio's creativity. As Engeström (2001) explains, the source of creativity is not inside a person's head, but it emerges in the interaction between a person's thoughts and his socio-cultural context. In design studios, communication and coordination between co-designers depend as much on different visual and physical aspects as they do on verbal aspects. During a typical collaborative design session, the type of information that is communicated between designers is multimodal, ubiquitous and touches the artistic, emotional and experiential side of the designers' thinking, in addition to their instrumental and practical reasoning.

Building on our previous work (Vyas et al. 2008; Vyas et al. 2009 and Vyas 2009), in this paper we focus on understanding collaborative approaches utilized by designers to aid creative support for ongoing design projects. We studied two industrial design departments over a period of eight months and explored three broad themes of collaborative practices. These are 1) externalization, 2) use of physical space, and 3) use of body. The *externalization* theme encompasses any kind of design knowledge represented onto three-dimensional, physical medium (e.g. sketches, models, prototypes) that can be used for establishing common-ground amongst co-designers. The *use of physical space* theme refers to a kind of ecological setup within a design studio, full of different types of design materials and artefacts (e.g. sketches, posters, timetables, to-do lists) that help co-designers organize, coordinate and manage their design work. The *use of body* theme refers to a collection of design practices where designers' bodies play an important role in exploring and communicating design knowledge with a group of co-designers. These broader themes encompass both pragmatic and instrumental factors related to design activities as well as inspirational factors that are important to aid creativity in the design profession. These themes are not mutually exclusive; on the contrary, their combinations are frequently used and they are frequently complemented by the other generic ways of communicating, such as, talking, overhearing and so on. Depending on designers' points of view, the rationale behind applying these collaborative practices range from clearly defining design problems, exploring new possibilities, easing communicative difficulties, to developing a communication language with co-workers.

Our motivation to do this kind of research is multifaceted. First, although research in HCI and CSCW has increasingly started focusing on the ‘design’ of interactive and collaborative technologies, ‘design as a profession’ is largely untouched as a subject of empirical study, with a few exceptions such as (Jaccuci and Wagner, 2003; Schmidt and Wagner, 2002; Robertson, 1997). However, we do acknowledge that there has been a sufficient amount of work done in developing tools and techniques to support design (Arias et al. 2000; Everitt et al. 2003; Hartmann et al. 2006; Maldonado et al. 2006). Secondly, as a part of creative industry, design cannot be easily formalized or rationalized to a specific set of activities, tasks or other kind of stereotypes. For example, traditional ways of communicating and collaborating may not be so important for the design profession (as we will see later). Hence, there is a need to understand how designers differ from other knowledge workers in terms of their working practices. Thirdly, we believe that in order to better support designers’ work and to develop new collaborative technologies, we need to understand how collaborative practices of designers enable creativity in their everyday work. An empirical investigation is required that specifically looks into the ubiquitous, collaborative and material nature of design practices.

In the rest of the paper, first, we will briefly describe background work that signifies the importance of embodiment in design work and some examples of augmented design environments. Next, we will describe our approach and methods used in understanding design environments in two industrial design departments. Next, we will describe the results of our study, focusing on the three themes of creative collaborative practices. And in the last section, we will discuss the implications of our results and provide a conceptual vision for developing technology to support collaborative design.

Related Literature

Our everyday communications and coordination acts go beyond linguistic signals and involve the use of material artefacts, locations and physical spaces (Clark, 2005). In fact, CSCW studies have increasingly shown the importance of material artefacts in coordinating distributed and co-located work (Hutchins 1995; Schmidt and Wagner, 2002; Sellen and Harper, 2002). Several authors (e.g., Kidd, 1995; Kirsh, 1995; Vyas, 2009) discuss how individuals intelligently make use of physical space and its affordances, in order to establish communication within a group. Advocating the use of ethnographic studies for designing systems, Randall et al. (2007) indicate three major aspects of artefacts that are relevant for understanding group work: ecological, coordinative and organizational aspects.

Amongst the empirical work on understanding design practices, Tang’s (1991) classic study focuses specifically on collaborative drawing, using observational video-tapes of three to four people collaborating at a table. Tang identifies several

features of collaborative work activity that should be taken into account when designing collaborative technologies. These are: 1) the importance of gestures, 2) drawing space as a resource for collaboration, 3) the importance of the process of collaborative drawing itself (instead of the final result), 4) recognizing the mix of simultaneous activities, and 5) the spatial orientation of collaborative workers. Jacucci and Wagner (2003) study the everyday practices of students at an architecture design laboratory. Their focus is on integrating ubiquitous computing technologies to support students' embodied interaction and contextualize these technologies to architectural design situations. Their ethnographic research shows the importance of material richness and diversity of material artefacts. They also register the distributed character of architecture learning and the use of space as a resource for collaborative interactions. The coordinative nature and the resourceful materiality of informational artefacts such as architectural maps or physical models are echoed by the work of Schmidt and Wagner (2002). In their later work, Jacucci and Wagner (2007) show how the materiality of informational artefacts plays an important role for creativity.

Hornecker (2002) uses an experimental setup where a group of co-located participants uses an assembly of three-dimensional objects in order to carry out paper prototyping as a design activity. Generating implications from a set of video recorded paper-prototyping sessions, her goal is to develop a graspable interface using table-top display technologies in order to support co-located design work. She focuses on the role of embodied actions such as use of gestures, parallel activities of participants and alignment of gestures with design artefacts and talks. A similar study is done by Robertson (1997), who develops a taxonomy of embodied actions of designers while working on cooperative design projects. She suggests that the public availability of different artefacts and embodied actions of distributed participants in a cooperative process could support communicative functions. She also argues that flexible and mobile access to the publicly visible information could improve coordination.

On the technological advancements in supporting design activities, we observe that researchers have focused on supporting embodied interaction in their technologies utilizing tangible and ubiquitous computing. Envisionment and Discovery Collaboratory (EDC) is one such platform that integrates two working spaces where stakeholders can incrementally create a shared understanding through collaborative design (Arias et al. 2000). The Distributed Designers' Outpost (Everitt et al. 2003) is a remote collaborative system that allows designers to use physical post-it notes to support discussion while designing websites. The application allows synchronous communication between distant designers through the use of 'transient ink' and 'remote shadow' mechanisms in order to coordinate design tasks.

Maldonado et al. (2006) developed the iDeas design ecology, a collection of tools that combines a browser for text and sketch-based design content, mobile

input mechanism for field observation data, and a vertical surface for collaborative creation and presentation. The d.tools toolkit (Hartmann and Klemmer, 2006) supports iterative prototyping of information appliances by allowing integration of design, test and analysis activities.

Understanding Collaborative Design – Our Approach

We investigated collaborative design practices in two industrial design departments in academic settings. Our ethnographic approach was informed by ethnomethodology (Randall et al. 2007). We intended to understand the everyday work practices of designers, methods and procedures they use to support their work and the resources they use to make sense of their design world. We used naturalistic observations, contextual interviews and video recorded collaborative design sessions of designers and design students. Our fieldwork lasted approximately eight months.

In the naturalistic observations, we studied the collaborative aspects of the design studios. Our goal here was to understand the natural circumstances of designers' collaboration, the tools and methods they use, and how the creative process of design is achieved. We had contextual interviews with 10 Master's students of industrial design and 5 designers / design researchers. We asked questions on individual ways of designing and on how designers understood creative ways of working. We asked how they brainstorm, what methods they use to come up with design concept, how they convey ideas to each other, their preferred tools for designing, the perceived advantages of using such tools, and so on. We took opportunities to record design sessions of groups of student designers. In some cases, we were participant observers collaborating with design students and recording their design proceedings.

In our analysis we identified three major themes of collaborative practices amongst the designers: *externalization*, *use of physical space*, and *use of body*. Our aim here is to show how creativity becomes an integral part of designers' work when they apply these collaborative practices while working in groups. We also want to stress that these collaborative practices are not used separately in all the cases and are often used in combination with each other.

Themes of Creative Collaborative Practices

In the following, we give specific attention to the three themes of collaborative practices of designers: 1) externalization, 2) use of physical space and 3) use of body. Figure 1 shows a typical scenario of a design session, where all three themes of collaborative practice play their parts.



Figure 1: A typical collaborative design session at an industrial design department.

Externalization

This theme was frequently observed as a major resource for establishing and enhancing creativity as a collaborative process. Externalization carries a broad range of design practices and activities: externalization of thoughts, of ideas and of concepts on a range of physical media. Artefacts such as paper sketches, drawings, posters, cardboard, clay or foam-models, and physical prototypes are examples of design externalization. Designers' externalizing practices vary over time (at different stages of design), in modality (from paper sketches to physical models), in purpose (exploratory or definitive), and are subject to individual preferences. In a single design project, design practitioners produce and use a plethora of design artefacts to support their work. These are constructed and used in and through an ongoing process of design. Within the context of industrial design, the externalization theme can be seen as a 'mediator' as well as a 'product' of cooperative design.

CSCW studies have shown that artefacts such as papers play a critical role in supporting social interaction and collaboration (Sellen and Harper, 2002). For designers, paper-based sketches have also shown coordinative advantages (Baskinger, 2008). With examples from our fieldwork we will describe how externalization plays a collaborative role in different activities and aspects of design: exploration, thinking by doing, coordination, and empathy and experience.

Exploration. Designers explore new ideas and concepts at various stages of their design cycle using different material artefacts such as sketches, mock-ups, models, and working prototypes. The goal here is to spend reasonable effort in order to get a partial result quickly. As one designer commented, "*in order to make design decisions you need to do explorations and for that you need to make*

different levels of prototypes". In design, everyday externalization practices involving sketches, foam or card-board could help designers explore new design ideas without too much effort. These types of external representations help designers to establish a creative sensibility. For example, sometimes sketching is used for visualizing designers' thinking as it stimulates creativity not only within their head but also with their hands. Figure 2a shows a brainstorming session where a group of designers are externalizing their ideas on post-it notes and at the same time giving a formal structure and category to their material. As one designer commented, *"Sometimes it is also useful to get something out of your head (externalize the ideas). When I have a lot of ideas and I know that some of these are not good, I just try to make a sketch of all of them and so that even some less important ideas are stored somewhere. I think it's a good thing that it gets me going."*




Figure 2: A structured brainstorming session using post-its to explore new ideas (a). Explorations of the effect of combining smoke and light (b).

We also observed that there are things that designers cannot easily envision through only drawing or sketching. They have to practically apply their ideas through different forms and textures of design models and prototypes to get a feel of their products. This kind of physical model allows designers to extend their mental conceptualization of their product to a sensory one. Figure 2b is an example of exploring the effect of smoke and different light colors in different shapes of glass. The idea here is to explore which combination would be suitable for a given situation. This designer explains that *"there are certain things that you cannot envision in a normal situation, things like "smoke". So in order to understand the behavior and interaction with smoke and utilizing it into design you have to build some things and play with it."* By joining the exploration of smoke with different kinds of lights, the designer explains, *"even by playing with a light I can get several ideas about new ways of interacting with lights, like blinking, fading, making patterns, so expressing new behaviors through the use of lights and different colors of lights. This opens up my visualization skills and*

provides new spaces for design. In this case if I just sketch this smoke with light, I wouldn't get that feeling. Here you can play with your hands, move the smoke around, this is a very different kind of design expression and gives me a different feeling."

Thinking through Doing. Designers communicate through a varied set of design representations often involving different materials, modalities and scale. To an extent, the whole design practice progresses through the use and manipulation of these representations and iterative refinements of both the conceptual and physical forms of products to be designed. Through externalization designers can visualize their ideas and concepts by actually creating them (putting things into practice) and not just by thinking about them. The physical activities and tasks that designers carry out allow them to think about the design of their products in a better way. During an iterative design process design artefacts such as sketches or models 'talk back' to designers (Schön, 1983). The epistemic knowledge developed during the process of constructing different design artefacts and externalizing design ideas leverages the way designers deal with elements of surprise and unexpectedness.



Figure 3: Externalizing design knowledge on different materials such as paper based sketches (a) and physical models using clay, foam, cardboard and plastic (b). (Photo: courtesy of Connie Golsteijn)

Our fieldwork on designers underscores the centrality of ‘thinking through doing’ (or thinking through externalizing). It was observed that a single design team would collectively develop an average of 50 to 100 external representations of their design ideas, depending on the project. These vary from paper based sketches or cardboard models to physical models. Because different styles and levels of fidelity of a representation yield different perspectives, meanings and experiences, externalizing ideas through a variety of prototypes affords a richer understanding of a design. Figure 3 shows two different examples where different design representations are used to support discussions. Figure 3a shows a design group using a collection of paper based sketches, whereas figure 3b shows a table

full of physical models made of clay, foam, cardboard and plastic. Being able to create more than one representation and alternatives of an idea and to try them out is in fact a major requirement for supporting creativity (Fischer, 2004). The thinking though doing theme suggests that the effort invested in developing different design alternatives helps co-designers to compare and judge important aspects such as the difficulty of building the final product.

Coordination. Several CSCW studies have shown that material artefacts play an important role in coordinating co-located and distributed activities (e.g. Sellen and Harper, 2002; Hutchins, 1995). Externalization of design ideas supports coordination within a team. The materiality of design artefacts provides information about the way they are created, used and manipulated, as well as about the process of design. Importantly, the temporality serves not only as indicative of different stages of a design process, it also serves accountability (planning, managing, budgeting, and so on) during the design work.




Figure 4: Group discussion of set of sketches (a), and result of a brainstorming session (b).

Different externalization techniques lead to creative methods of communication within a design team. Externalizations support creativity as they provide opportunities for others to interact with, react to, negotiate around, and build upon an idea. Externalizations contribute to a common language of understanding amongst a group of designers. For example, figure 4a shows a group of designers discussing different sketches at a table. Figure 4b shows results of another brainstorming session where the cooperative nature of design artefacts helped to develop new alternative concepts. The important issue here is that the materiality of different design representations can afford and trigger different collaborative actions in the team.

Empathy & Experience. In our field study, we saw several examples where designers created design representations based on observations of the real users. They tried to provide as much empathy towards the users through the development of such representations. One of the most powerful human capabilities relevant to designers is the intimate incorporation of an artefact into

physical experiences to the point where people perceive that artefact as an extension of themselves; they act through it rather than on it (Klemmer et al. 2006). Additionally, different design materials and artefacts allow direct and bodily engagement and hence broaden communicative resources by evoking sensual experiences. The multi-modality and ability to support and convey information through all senses, makes the use of a design artefact experientially rich (Vyas et al. 2009). In the case of joint design activities, co-workers do not just interact with these artefacts when they are designing, they actually get the feeling and experience each other's activities through these artefacts. The communication channels that are established by these multi-modal artefacts go beyond facilitating basic task-oriented activities.

Use of Physical Space

This theme refers to how design practitioners utilize their physical surroundings within a design studio in order to support collaboration and creativity in their work. In both of the design studios that we studied, we saw design teams use their office walls, whiteboards, clipboards, wooden panels and so on as carriers of their design-related information. The types of information that are attached to these spatial objects have instrumental and productivity related functions and can be seen in the form of design ideas, sketches, to-do lists, project-related information, work-in-progress data and other organizational details. At the same time, they also carry inspirational, provocative and other non-instrumental details such as posters and innovative design sketches. The way information is represented in the space provides indication about collaborative and methodic practices of designers (Vyas 2009).



Figure 5: An example of creative ecology in a design studio.

Figure 5 gives a glimpse of a section of a design studio where a design team has used clipboards, large card boards and movable tables to develop a creative environment. In addition, there is information about project plan, post-it notes, design sketches on the clipboard, as well as the prototype on the table. An environment such as this establishes a ‘creative ecology’ within a design studio both at personal and social level. In the following, we will discuss how arrangements such as these help in establishing creativity.

Elaborate the Problem. One of the reasons to utilize space in such a way is to elaborate and divide design challenges so that detailed descriptions of different aspects of design can be generated, which in turn would help in resolving a particular situation. The way physical space allows the representation of design tasks can affect designers’ reasoning abilities and performance. As one designer suggested, *“I normally try to visualize all the material and data that I collected from my user studies and try to find out patterns and explore design opportunities from this data. I then make my own sketches and models and keep all these in a way that can help me find out new ideas”*.



Figure 6: A shared design environment, with pictures of different field studies and observations on the walls and desk (a). Detailed personas on a wall of a design studio (b).

Several examples of this were seen in both of the design studios. Designers keep, for example, pictures from ethnographic or other field studies on their office walls and around their desks (figure 6a), or develop persona archetypes of their potential user groups and stick them on their shared working spaces (figure 6b). The aim here is not just to solve a design problem but to collect greater and useful insights into a given situation so that solutions can be envisioned.

Awareness. Within an ongoing design project, designers deal with a plethora of design materials, and being aware of different ‘happenings’ is an important issue. We observed that the way designers keep project-related design materials on different spatial objects within their studios improves the visibility and provides an overview of the work being carried out. Understanding how design artefacts within a work environment are organized, configured, manipulated and

handled supports the awareness of co-workers' activities and, hence, contributes to the coordination of work. Design iterations, methods, and conventions can be easily extracted when design artefacts and related materials are kept in public visibility using physical space. The visibility of design activities is also manifested in and through the use of these artefacts. At the same time such a creative space could provide opportunities to reflect on the ongoing project and to allow designers to change, combine or divert aspects of their design process.

Organize and Manage. Design being a collaborative process requires organizing and managing the work of co-designers. The spatial aspects within design studios also play a role in supporting the organization and management of design projects. Figure 7 shows two examples (a & b) where design teams have used clipboards and movable drawing boards to show information related to project plans, data generated from brainstorming sessions, design concepts, work division within a team and to-do lists. Each individual piece of design-related information has a strong, even explicit link to some aspect of the project at hand. The ecology of these pieces information creates an information rich environment needed to stimulate creativity and to develop novel ideas.

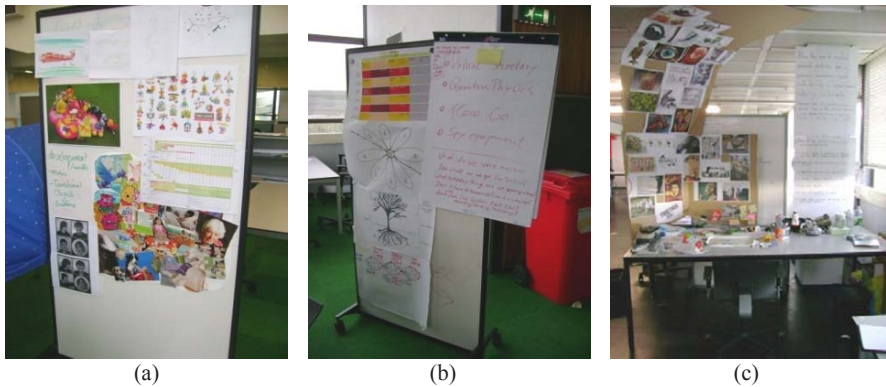


Figure 7: Shared clipboards full of design-related materials to organize and manage ongoing projects (a & b). The personal workspace of a designer (c).

Personal vs. Shared. The way different information and design artefacts are arranged within design studios establishes a vague distinction between personal and shared spaces. The above figures 7a and 7b show a physical space that is shared by a group of designers. However, designers also have their individual working space that they organize based on their own personality and reasons. As one of the designers commented about his private space, *“the space allows me to organize my work and get reminded what I am doing daily. Also for the purpose of communicating with my peers I can very easily show what I am doing.”* As can be seen in figure 7c, these artefacts are indicative of different phases of the design process, the current state, and future planning. Another designer commented,

“depending on the phase of the project, I arrange my surroundings. It’s important for me to have these artefacts around so that I can register where I am at in the project”. Hence, these design artefacts were markers for reminding. Personal spaces also allowed designers to create a portfolio-like arrangement of their workspaces expressing an identity or self-image.

Use of Body

During ongoing design projects, designers accomplish activities and tasks not only through their internal cognitive processes but by utilizing cooperative ‘embodied’ actions (Robertson 1997). The third theme that we discuss here is about how the specific use of designers’ bodies helps in establishing creativity in collaborative design practices. The use of the body theme is central to externalization and utilizing the space (the above two themes) in all design activities. Designers creatively make use of their bodies while talking, while explaining a design sketch or in referring to spatial arrangements within a design studio. While the use of gestures and other bodily representations for discussing design ideas is common in design studios, there is an increasing use of design methods such as role playing, body storming or design choreography in groups (Hummels et al. 2007). Using these methods, designers explore and experience design possibilities for themselves, intentionally make these ideas public and allow other designers to reflect on these ideas. Here the design cooperation is achieved by the mutual perception of these actions as the basis for the ongoing creation of shared meanings in a particular design task. The use of bodies can be seen in different design stages to support different needs. In the following we will explain how the use of bodies helps in creativity.



Figure 8: Exploring design possibilities through performances. (Photo: courtesy of Rob Tieben)

Exploring Interactive Concepts. It has been suggested that bodily movements are suitable as a design technique, as our bodies convey emotions as well as geometry and interactions (Hummels et al. 2007). Role play methods allow designers to imagine and empathize a given design challenge. A physical activity is a primary source here to explore new possibilities. In our fieldwork we found that many of these bodily actions were aimed at better understanding of the design task context and at exploring new possibilities. Figure 8 shows two examples of exploring design possibilities. Here, the participants, using different bodily patterns, are exploring the possible behaviors of the product to be designed. The vividness of these experiences and the bodily understanding of a given design situation help designers to make better design decisions (Buchenau and Fulton Suri, 2000).

Improve Communications. Our verbal languages may not be enough when communicating issues related to complex technologies. While designing new technologies or products, designers have to think about out-of-the-box ideas that may be difficult to articulate using verbal means. One of the main objectives of applying role play methods is to communicate early design ideas and concepts in an engaging and participative way that could establish common-ground for the group of designers (Buchenau and Fulton Suri, 2000). Additionally, many product designers need to deal with issues such as branding, marketing and advertising. Methods such as role play help in dealing with all these issues in one package – that requires a combination of functionality, expression and communication.

Studies have shown that gestures, in addition to their purely communicative role, help lighten cognitive load when a speaker or performer uses them in combination with speech (Tang, 1991). Through role playing, a performer’s ability to map his/her actions to certain features or tasks of design could help in understanding the envisioned product.

Exploring new Experiences. Supporting appropriate user experience is amongst the main goals within the design profession (McCarthy and Wright, 2005). Our physical bodies play a central role in shaping human experience in the world, in understanding of the world, and in interaction with the world (Klemmer et al. 2006). In addition to exploring new ideas and improving communication possibilities, we also observed that the use of role play and other participatory methods provided new perspectives on bodily experiences. When designers enact a particular scenario, they go through a set of emotional and experiential “phases” that not only make their actions personally meaningful but also lead them to envision how a potential experience should be.

Improving Design Practices. Echoing the claims of Fischer (2004), we observed that being able to move around the design environment and to interact with different design-related artefacts and with other designers can help in the understanding and learning of creative designing. This was in fact an important rule-of-thumb in one of the design studios that we visited. One of the professors

of the industrial design department frequently advised designers working in the studio to “*move around and don’t just sit at the desks*” to generate creative ideas.



Figure 9: Design students collaboratively sketching – influencing and inspiring each other.

During the interview sessions with professional designers, we learned that on several occasions designers brainstormed by simultaneously drawing quick sketches and doodles on large sheets of paper in order to generate quick design ideas. Figure 9 shows design students at an industrial design department collaboratively exploring new ideas on a large sheet of paper. In close proximity, designers can influence and inspire each other and at the same time adapt to each other’s sketching styles. This theme suggests that creativity is an applied phenomenon, in full, creativity can be established by practicing and doing things in the real world, where bodies play a critical role.

Discussion and Design Concepts

The ethnomethodological approach allowed us to understand the current practices of designers to support creativity in their ongoing design work. In particular, the examples that are discussed in this paper point to the critical role of ‘material collaboration’ in supporting and enhancing creativity in the context of cooperative design. The three themes related to collaborative practices of designers that we have discussed here, namely, externalizing, use of space and use of bodies, provide insights into how material and physical signals can trigger creative thinking. We believe that there are important implications for the development of collaborative technologies for supporting professional designers. In the following we will describe these implications.

Spatial flexibility is an important factor for supporting group creativity of designers. It was apparent in our examples that designers develop a multitude of

design artefacts in the form of paper sketches, drawings, physical models and so on. The way designers keep these artefacts and organized them in their workspace affects their work organization, communication and coordination practices. It is this spatial flexibility of, for example, sticking sketches and drawings on a shared office wall or keeping physical models of different materials on a table that allows designers to discuss, criticize and explore new possibilities of their design work. In order to provide technological support for spatial flexibility, we need to think beyond desktop computers. Jaccuci and Wagner (2003) made an attempt to support spatial flexibility via mixing real work objects with virtual ones to support learning and collaborating amongst students of architecture.

Archiving materials used and produced during design processes helps co-designers get back to them whenever they need. There is creative value in allowing designers to associate and connect different design artefacts. We observed in our fieldwork that designers attach paper based sketches, drawings and posters to their vertical surfaces for different purposes, creating a technological environment that allows designers to archive these design materials in such a way that could lead to supporting creative thinking.

Encouraging movement is an important aspect for aiding designers' collaborative creativity. As it was seen in the examples, designers' physical movements during explorative design stages and while using methods such as role playing or body-storming, support creativity in group sessions. Supporting the call for Design Movement (Hummels et al. 2007), we believe that technology should not hinder designers' physical capabilities but, on the contrary, should encourage freedom.

Sustaining ubiquity of design practices, especially when people collaborate from remote locations, could be a challenging task for developers. As was observed in our fieldwork, information related to a design project can be found in physical space and in material artefacts, as well as in designers' ability to utilize these material aspects. We believe that technologies that support live transmission of audio-video links may be able to support designers' conversations, but the pervasive nature of design practices requires the creation of technologies that go beyond these conversational paradigms.

Supporting thick practices of designers is a design challenge that should be taken into account. By this we mean that any new technology should acknowledge and take into account the primacy of real-world design practices. Technology should not just bring new ways of working but instead improve flexibility for the designers to use their methods. In this sense, a technology should carefully integrate physical and digital worlds to enable the improvisation of practices that the real world could offer.

Keeping these implications as a base, we have developed two conceptual systems that could potentially be used to support collaborative design activities. These are 1) the resource sharing concept and 2) the live discussions concept.

Resource Sharing Concept




Figure 10: Resource Sharing concept on a table-top workspace

The first of our concepts, Resource Sharing (figure 10), allows creative collaboration between designers in a co-located situation. The Resource Sharing concept uses a tabletop interface which allows designers to discuss and share design resources related to their products or prototypes. The tabletop interface can generate the design history of a physical product once it is kept on the surface of the table. The table shows the digital versions of the product ideas, associated sketches, annotated drawings and other historically important details in a hierarchical format. The table supports the use of multiple physical products or prototypes. As can be seen in the figure 10, using this tabletop interface, designers can look back in time, re-view the options they considered and reflect on them. The tabletop interface also allows designers to make new sketches on the interface based on what they are currently discussing in a design session.

This concept uses tabletop technology to allow designers a kind of spatial flexibility compared to a typical desktop based system. This spatial flexibility allows designers to collaboratively access multiple design artefacts (e.g. sketches) at the same time carry out brainstorm activities. In a sense, the table-top interface provides a mixed-reality interface to discuss real-world objects and associated digital artefacts, and it allows designers to sketch new design ideas on the surface of the table. The ability to connect, associate and compare multiple design artefacts on the table surface could enable designers' creative brainstorming activities. The tabletop interface does not impose any substantially new practice, it just allows new ways of interacting and storing design ideas in the table.

Realizing this kind of technology may not be too difficult as existing tabletop technology such as Microsoft Surface or Philips Entertaible can be used.

Live Discussions Concept




Figure 11: Live Discussions concept

The second concept, Live Discussions (figure 11), focuses mainly on remote collaboration, allowing designers to discuss three-dimensional and physical objects or prototypes as well as two-dimensional paper-based sketches without losing information. As can be seen in the figure 11, design studio A has a table with dedicated planes (surface spaces) to allow communication of different types of design artefacts. Design artefacts have a RFID tag attached and different planes on the table are equipped with RFID readers. Design studio B is located at a distant place with other members of the same team. Studio B has a large-screen touch display where the view of the table in studio A is shown, with the help of RFID tags and readers. The dedicated planes on the table help to adequately represent the two-dimensional and three-dimensional information. In Studio B, designers can point, annotate or draw on a particular part of a design object and simultaneously communicate via microphones. Studio B is equipped with a hi-resolution camera that shows the live feeds of Studio B onto a display located in Studio-A. The concept is partly based on the work of Everitt et al. (2003), where design brainstorming was made possible through the use of post-it notes.

This concept is based on implications from our fieldwork. By allowing spatial flexibility though the use of a dedicated design table we could allow to discuss both two-dimensional and three-dimensional objects and to brainstorm over a distance.

Conclusions

The observations and ideas discussed above do not address the entire range of practices of the design studio culture. The three themes of collaborative practices that we discussed cover a broad spectrum of techniques that designers use to aid creativity in cooperative design. Clearly, creativity is a critical aspect of design and needs to be supported through technological means. What has been presented here is an account of how creativity is applied by the designers of the two industrial design departments. An account of real-world design practices such as this could be very fruitful when we are to design collaborative technologies.

This study reflects the embodied nature of design practices. Our work shows: 1) how different externalization techniques utilizing seemingly mundane and simple design artefacts such as sketches, post-it notes, and physical models within a design studio play a role in supporting designers' everyday creative work; 2) how the intelligent use of physical space of a design studio helps designers to think creatively about their design work; and, 3) how bodies of designers play a pivotal role in experiencing and envisioning design aspects. The rationale behind applying these collaborative practices ranged from clearly defining design problems, exploring new possibilities, easing communicative difficulties, to developing a communication language with co-workers.

Acknowledgement: This work is supported by the EU Project AMIDA (FP6-0033812).

References

- Ackerman, M. S. and Halverson, C. (1999): 'Organizational Memory: Processes, Boundary Objects, and Trajectories'. In *Proc. of the 32nd Annual Hawaii international Conference on System Sciences*- Volume 1. HICSS. IEEE Computer Society, Washington, DC, 1067.
- Arias, E., Eden, H., Fischer, G., Gorman, A., and Scharff, E. (2000): 'Transcending the individual human mind—creating shared understanding through collaborative design'. *ACM Trans. Comput.-Hum. Interact.* 7, 1 (Mar. 2000), 84-113.
- Baskinger, M. (2008): 'Pencils before pixels: a primer in hand-generated sketching'. *Interactions*, 15, 2 (Mar. 2008), 28-36.
- Boess, S. U. (2008): 'First steps in role playing'. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*. (CHI'08) ACM, New York, NY, 2017-2024.
- Buchenau, M. and Fulton Suri, J. (2000): 'Experience Prototyping'. In *Proc. of Fourth Conference of Designing Interactive Systems (DIS'00)* ACM Press: NY, 424-433.
- Clark, H. H. (2005): 'Coordinating with each other in a material world'. *Discourse and Society*, 2005, 7: 507-525.
- Engeström, Y. (2001): 'Expansive Learning at Work: Toward an Activity Theoretical Reconceptualization', *Journal of Education and Work* 14(1): 133-156.
- Fischer, G. (2004): 'Social Creativity: Turning Barriers into Opportunities for Collaborative Design.' In F. deCindio, & D. Schuler (Eds.), In *Proc. of the 8th Conference on Participatory Design Conference (PDC'04)*, Volume 1, ACM Press, NY, 152-161.

- Hartmann, B., Klemmer, S. R., et al. (2006): 'Reflective physical prototyping through integrated design, test, and analysis.' In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology* (UIST '06). ACM, New York, NY, 299-308.
- Hornecker, E. (2002): 'Understanding the Benefits of Graspable Interfaces for Cooperative Use'. In *Proc. of 5th International Conference on the Design of Cooperative Systems* (COOP'02). IOS Press, Amsterdam, 71-87.
- Hummels, C., Overbeeke, K. C., and Klooster, S. (2007): 'Move to get moved: a search for methods, tools and knowledge to design for expressive and rich movement-based interaction'. *Personal Ubiquitous Comput.* 11, 8 (Dec. 2007), 677-690.
- Hutchins, E. (1995) *Cognition in the wild*. MIT Press, Cambridge, USA.
- Jacucci, G. and Wagner, I. (2003): 'Supporting Collaboration Ubiquitously: An Augmented Learning Environment for Design Students'. In *Proc of 8th European Conference on Computer Supported Co-operative Work* (ECSCW'03), Kluwer Academic Publishers, 139-158.
- Jacucci, G. and Wagner, I. (2007): 'Performative roles of materiality for collective creativity'. In *Proc. of the 6th Conference on Creativity & Cognition* (C&C '07) ACM Press, NY, 73-82.
- Kidd, A. (1994): 'Marks are on the Knowledge Worker'. In *Proc. of CHI'94*. ACM: NY, 186-191.
- Kirsh, D. (1995): 'The intelligent use of space', *Artificial Intelligence*, v.73 n.1-2, 31-68.
- Klemmer, S. R., Hartmann, B., and Takayama, L. (2006): 'How bodies matter: five themes for interaction design'. In *Proc. of the 6th Conference on Designing interactive Systems* (DIS'06). ACM, New York, NY, 140-149.
- Klooster, S. and Overbeeke, C. (2005): 'Designing products as an integral part of choreography of interaction: the product's form as an integral part of movement'. In *Proc. of 1st European workshop on Design and Semantics of Form and Movement*. New Castle, UK, 23-35.
- Maldonado, H., Lee, B., and Klemmer, S. (2006): 'Technology for design education: a case study.' In *CHI '06 Extended Abstracts on Human Factors in Computing Systems*. (CHI'06) ACM, New York, NY, 1067-1072.
- McCarthy, J. and Wright, P. (2004): *Technology as Experience*. MIT Press, Cambridge, MA.
- Randall, D., Harper, H. and Rouncefield, M. (2007): *Fieldwork for Design – Theory and Practice*. CSCW series, Springer-Verlag London.
- Robertson, T. (1997): 'Cooperative work and lived cognition: a taxonomy of embodied actions'. In *Proc. of the 5th Conference on European Conference on Computer-Supported Cooperative Work*, (ECSCW' 97). Kluwer Academic Publishers, Norwell, MA, 205-220.
- Schön, D. (1983): *The reflective practitioner: How professionals think in action*. NY: Basic Books
- Schmidt, K. and Wagner, I. (2002): 'Coordinative artefacts in architectural practice'. In M. Blay-Fornarino et al. (eds.): In *Proc. of the 5th International Conference on the Design of Cooperative Systems* (COOP'02), IOS Press, Amsterdam, 257-274.
- Sellen, A. and Harper, R. (2002): *The Myth of the Paperless Offices*. MIT Press, MA.
- Tang, J. (1991): 'Findings from Observational Studies of Collaborative Work', *International Journal of Man-Machine Studies*, 34, Elsevier, 143-160.
- Vyas, D., Heylen, D., and Nijholt, A. (2008): Physicality and Cooperative Design. In *Proc. of 5th Joint Workshop on Machine Learning and Multimodal Interaction*. (MLMI' 08), Lecture Notes in Computer Science, Springer-Verlag. 325-337.
- Vyas, D. (2009): Artful Surfaces in Design Practices. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*. ACM, New York, NY, 2691-2694.
- Vyas, D., Heylen, D., Nijholt, A. and van der Veer, G. (2009): Experiential Role of Artefacts in Cooperative Design. In *Proc. of 4th International Conference on Communities and Technologies* (C&T' 09). Penn State University, PA, USA. ACM Press: NY.

‘Talking about (my) Generation’: Creativity, Practice, Technology & Talk

Authors: David Martin¹, Jacki O’Neill¹ and Dave Randall²

¹Xerox Research Centre Europe, Grenoble, France

²Manchester Metropolitan University, Manchester, UK

¹{firstname.surname}@xrce.xerox.com, ²d.randall@mmu.ac.uk

Abstract. This paper describes the findings of an ethnomethodological enquiry into the work of graphic designers. We explore the collaborative nature of graphic design as undertaken by a small team of designers working in a packaging design company. In doing so, we attempt to explicate the way in which practice, talk and technology are intricately bound up in such a way as to constitute a creative process. We describe a series of scenic features, ‘orderings’, and ‘talkaboutables’ which are characteristic of this process and which may be entailed in other creative contexts and hence can be important topics for CSCW design for creativity.

Introduction

"I was recovering from a weeklong stint of design work in my Chicago studio. In those days I was in the habit of bolting my studio door and immersing myself in the heady universe of packaging design. It was my true creative calling ..."

David Marusek, 'Counting Heads'

"In order to get clear about aesthetic words you have to describe ways of living. We think we have to talk about aesthetic judgments like ‘This is beautiful’, but we find that if we have to talk about aesthetic judgments we don’t find these words at all, but a word used something like a gesture, accompanying a complicated activity"

Wittgenstein 1966, 11.

This paper describes the findings of an ethnomethodological ethnography of graphic designers. We explore the collaborative nature of graphic design as undertaken by a small team of designers working in a packaging design company. We show how the creative process is intricately bound up with talk about the work, orientation to various 'known in common' background features, and shaped by the various technologies used, showing through a series of examples the social and methodical features of the creative work.

Creativity is an interesting topic for the CSCW community, since it has in the past often been conceptualized as a 'black box' - as entailing 'Eureka' moments (e.g. in major scientific discoveries) or as a solitary, perhaps cognitive, pursuit (e.g. music composition). Historical treatments show how our conception of the creative process has been progressively transformed and widened from in ancient Greece, where only poetry was considered creative (even visual arts were considered only to copy nature), through the transformation of the visual arts from the artisanal to the 'creative' in the Renaissance, towards a heroic conception of the creator from the 18th century.

More recent research has begun to explore how creativity is often a collaborative achievement (e.g. Fischer, 2005; Cook, 2005; Jacucci and Wagner, 2007) and certainly creative work is always mediated through technology (for recent discussions of materiality, new technology and the creative process see Jacucci and Wagner, 2007; Eales, 2005). Although such work has begun to explore the situated, technologically-mediated nature of creativity, there are still relatively few studies of creativity 'in the wild', as opposed to more self reflective approaches (e.g. Schneiderman, 2000) or the vast literature that seeks to model creativity (c.f. Greene, 2001), particularly given the broad application of the concept across many domains.

Creativity is now a term that can be applied to arts, sciences, business and design and there is a strong interest in developing technologies which can enhance creativity. Hence, in the world of 'design', there is increased interest in (and a progressive convergence between) the design of 'systems' of whatever kind and design seen as a 'creative' activity. This move has occurred as computer systems seek to support ever more varied and complex activities in novel and innovative ways. This convergence is arguably prompted by at least three distinct 'moves' in design-related arenas such as CSCW. Firstly, as we have 'moved out of the control room' (Hughes et al. 1994) then so have we recognised the heterogeneous nature of skills and expertise in complex environments, where, for example, systems need to support heterogeneous user constituencies across organisational boundaries, and hence the difficulty of representing them (see Ackermann et al, 2003) leading, *inter alia*, to the reconsideration of concepts such as 'awareness' (see Schmidt, 2002). Secondly, with the advent of Web 2.0 applications, some part of the project at least has involved embedding symbolic information into computer systems in such a way that semantic information can become 'machine

readable' and the handling of multi-media information in such a way that it can be tagged, taxonomised and indexed. Creativity and other social activities, for all sorts of users, are supported and 'catalogued'. Folksonomic applications, for instance, are now commonplace, and are evident in sites like Flickr and social bookmarking sites like Delicious. Furthermore, the Web 2.0 infrastructure allows the creative ad hoc development of new applications, for example, through mash-ups. Thirdly, CSCW itself has expanded into a range of domains that would twenty years ago have been unthought-of, including work on domestic life and on public computing (e.g. in the realm of games and arts) of one kind or another.

These factors and others, we believe, have led to something of a sea change with respect to the recognition of creativity in design. This is reflected in attempts to conceptualise 'creativity' in design (see Herrman, forthcoming) and to 'scope' design issues as heterogeneous and creative processes. Nevertheless, this is not without challenges of its own. New relationships between designer and user may be entailed, along with new problems of conceptualising exactly who our users might be and how they might be mobilised.

This is not to say, of course, that the use of technology in the creative process is new for it evidently (for anyone familiar with the paintbrush) is not. Indeed, art history has made various sporadic attempts to understand the relationship between technology, mundane practices and creativity (see for instance, Baxandall, 1974; Benjamin, 1982; Fischer, 1970), albeit to varying effect. What is relatively new, however, is the use of *collaborative technologies* to support the creative process in heterogeneous environments. Graphic design is, we argue, one such case. It is commonly thought of as a creative profession: it involves artistic abilities and is engaged with producing designed-for-purpose products which are realised on various materials and media (for example, paper, fabric, plastic and digital); the work of designers interlinks with other organisational processes, and involves the deployment of heterogeneous knowledge and resources, both technological and otherwise, not to mention socially shared methods and practices for using them, including language and discourse. All of these are implicated in creating, understanding and evaluating a developing design.

Method and field site

Ethnography has played an influential role in CSCW, providing as it does detailed insights into work (and other) practices. The ethnomethodologically-informed variety, employed in this study, has consistently identified and emphasised the detailed explication of the ordinary, practical ways in which people go about doing the things that they construe as relevant, important or necessary in the context which they inhabit and produce. Such ethnographies have progressively exposed the knowledgeable, artful ways in which participants orient to their work using and constituting that work through the use of technologies and

other artefacts (see e.g. Button and Sharrock, 1997; Goodwin, 1994; Whalen, 1995). Such work has sometimes been tightly coupled with design, sometimes more evaluative and increasingly has stood as an informational resource about practice. We would suggest that the latter instance, one recommended by, for instance, Dourish (2007), is peculiarly well-suited to problems of the 'creative'. This is because (as pointed out by Randall et al, 2007), as we move out not only from the control room but also out of the bounded organisation, decisions about who and what to study, and in what situations, become increasingly complex. This third use might be considered as apolitical insofar as the aim is to explicate practice rather than to directly influence design. Like Dourish we would contend that this form of ethnography is of specific relevance to CSCW in situations where we may have relatively little to go on - conceptually and empirically - in efforts to open the 'black box' of creativity, as it unfolds in practice.

The graphic design company we studied - 'Box Group' - is a small independent start-up agency, incubated in a University, which specialises in packaging design. Their work ranges from innovative packaging concept design for large UK retail customers to more mundane packaging (and other) graphic design work for small and medium businesses (SMBs). They have 4 full-time employees (a creative director, 2 graphic designers and an administrator) who are all located in the same small room, and 2 part-time designers. The creative director concentrates on innovative packaging, sometimes in concert with the part-time designers. The full-time designers do the more run-of-the-mill design work for SMBs, doing both the creative design (ideas and concepts) and the mechanical design (the final accurate product). This ethnography was conducted over two separate weeks a few months apart. The study was observational with ad hoc interviews conducted from time to time to clarify or understand elements of the work. Several hours of audio recordings were made each day. Some video was also collected, as well as photographs, many notes and various design materials. The material was analysed from an ethnomethodological perspective (e.g. see Randall et al. 2007).

Social and methodical features of creativity

In this section, we examine how the design context both inspires and constrains the design space and how the tools and technologies enable searching, comparing, choosing and detailed editing work. The designers in Box Group work with a range of technologies, applications and resources. The key technology for them is Adobe Illustrator© where the detailed design is produced. Illustrator is very good for producing graphical design elements but text and photographs can also be brought in, assembled and manipulated in a number of ways. The designers also use Adobe Photoshop© for the assembly and manipulation of photographs. In terms of other resources, the designers use various on-line photograph and graphics/icons libraries like Getty Images and font libraries like Dafont, some are

subscription and some are free. Customers may supply specific fonts and logo files to be used in their designs. They also use digital cameras, scanners, printers and traditional technologies like pens, pencils, paper, scissors and glue.

In what follows we characterise some of the activities involved in this kind of design work. We are not, it should be stated, suggesting that they take place in any determined order, for as with many other forms of work, we see a variety of iterations in play.

Assembling Resources

We have pointed to the way in which organisational studies have increasingly recognised the heterogeneous nature of organisational life. One aspect of this is, the inter-organisational - a common feature of which is 'customer facing' work. An over-arching 'scenic' feature of creative design work is that it is customer-facing. Graphic designers draw a number of distinctions about work for a customer. Most importantly, they distinguish concept design from 'artwork', or mechanical design. In the first, their work is looser, providing 'concepts', i.e. more or less 'accurately' realized sketches of ideas, containing the different textual and graphical/pictorial elements. In the second, they produce the finished product to specification. This is about correctness and detail, and may involve a series of products over a range. It is time consuming and meticulous, and the designers talk about it as the more humdrum part of the work. It is 'small' creativity to them as opposed to the freer ideas work in the earlier stage. However, creativity also depends on the freedom of the brief, the will of the customer for something 'out-there' and so forth. Some pieces of work are restrictive and require only an update, in a similar style, with similar fonts, colours etc. The opportunity for creativity is low. Even so, whilst the opportunity to be creative is desired, being given a free reign (i.e. with a customer unclear about what they want) is not. Clear ideas from the customer are very important in shaping which way to go.

Graphic designers, in other words, do not work in an unconstrained space, for they are providing a service for a customer and a first step in a new design project is to collect ideas, requirements and 'parameters' that allow them to circumscribe the design space – to fix elements of the design. Thus, when designers talk to a potential customer they are engaged in the process of 'constructing the brief' such that the initial brief is made tractable. This is done, in part, by assembling relevant information. Indeed, designers suggest that a large design space and freedom for creativity can be a dangerous thing; it can stop them getting on with the design and it can make it difficult to know if the customer will like what they have produced (N = Neil designer, E = ethnographer):

1. N – Well sometimes a free space can be even worse because sometimes its better with parameters, it's easier with restriction...cos it, speeds up the decision making process... it means you can't go down that sort of route, you can't go down blind alleys

2. E – And even time can be a restriction?
3. N – It is, because time is always money as we all know... and that always, in the back of your mind... sometimes that's a good thing too because you won't pfaff around, you won't sit with your feet up and stare out a window thinking it'll come, it'll come

Designers' orientation to various constraints, then, consists in their knowledge of and about a variety of relevant activities which can and does include:

- The receipt from the customer of a written brief, some examples of current products (the products themselves and the artwork for them), additional information from interviews and telephone conversations. They may even try out the products e.g. chocolate bars and sweets.
- The gathering of background research material – i.e. packages and designs of similar and contrasting products in the same general area.
- Knowledge of the 'library' or collection of materials distributed around the office that they have collected and assembled magpie-like; designers like picking up designs and other objects that appeal to them, and they do it as part of their work.
- Access to online resources such as libraries of fonts, images and so on.

These collected artefacts get used as context setting and shaping devices when the designers, often collaboratively, employ their knowledge and skill in relation to them when they develop and appraise their designs.

Determining Relevance

Of course, in principle any number of possible resources could be used, and in some way they have to be selected. An example of the collaborative work entailed in this is provided below. Here, the job in hand is the packaging re-design and re-branding carried out for two specific brands of sweets ('mini mix and mini fruit drop mix'). From the start we can see the collaborative nature of design work as one of the more experienced designers, Jim (J), relays information on the products, the written brief and his conversations with the customer's agent to a less experienced colleague, Annie (A) who will do the work:

1. J – ... So the thing they want to shout about is the fact that they are a ten calorie sweet ... and the name of them ... I'll give you the brief afterwards mini mix and mini fruit drop mix
2. A – They don't have any specification on fonts or anything
3. J – No we've got complete, whatever we want – the new design (**reading**) must advertise the low calorie count, advertise no artificial colours and preservatives, present that product as a more value looking brand than the current SweetCo packet so it can't be that it looks high value its got to look like a fun sweet ... but if he's saying it has to look like a lower value brand than that one you've got there then that looks like the cheapest brand you've ever seen



Figure 1: Old designs for the sweet packets




Figure 2: Annie's new design

Jim begins by saying that the *'thing they want to shout about'* and lists two things; the *name* and the fact that they are *ten calorie sweets*. That suggests a prominence for these two pieces of text on the packet. Annie asks whether there is a specification of fonts or anything else. Jim answers *'No, we've got complete, whatever we want'*. This does not translate into we can do whatever we like, rather a number of design constraints are clearly worked up in the description and conversation. Jim goes on to read the brief which specifies further textual requirements and a statement about the brand being a *'value looking brand'* in comparison with another company brand. What *'value looking'* might mean is further elaborated when Jim picks up the packet that he believes suggests *'higher quality'* to the customer – *'... that looks like the cheapest (i.e., most extreme 'value') brand you've ever seen'* – suggesting that the new design should not try and look cheaper than what the customer thinks of as higher quality¹. Instead (see turn 4 below), he suggests they should design something comparative to another popular UK brand – *Celebrations* – and different from another more expensive brand – *Thornton's*. This type of comparison is a common feature of design work – since it is collaborative work the designs need to be articulated (whether amongst the designers themselves or between the customer and designer) yet the design itself is a visual matter. So the designers use comparisons such as these – which might be thought of as product discourse, along with other *'design speak'* to inspire and articulate designs². They can be considered as a form of membership categorisation device [Sacks, 1995], employing members' knowledge of what particular product designs are meant to convey, and serve as starting points for design.

¹ This is also a joke about customer taste.

² Inspiration works in complex ways – it may be about *'categorising'* the product with other products, it may be general features, it may be just a detail, or sets of them, or a font *'type'*, a use of colours etc. How similarity and difference are to be realised for *'this'* product is creatively produced through design activities and in relation to talk like this.

Design Elicitation: The Product and the Range

As the conversation continues Jim suggests how Annie should proceed and makes some suggestions: draw a bag shape, have a window in it, try fonts and brands, less photographs and more graphics:

4. J – So ignore it but don't make it look like a premium brand ... make it look like a nice cheery bag, of, 'Celebrations' or something like that sort of you know what I mean, it doesn't look like Thorntons but it doesn't look like that, on the table ... we will produce a number of design concepts for both mixes and present for comment, that's all that we need to do at the moment. So all you need to do is ... a bag shape and mess around with a design, and you can have a window in it as I say ... they always do have a window in it ... now ... I bet that's the one that he thinks is higher class brand at the moment that, it doesn't have to look as premium as that so less photograph's and more, more graphics, but all I want is, really free, ... just messing around with fonts and brands and stuff like that...

Being product designers, the physical shape and material of the packaging is also subject to design, but is also used as an initial starting point, being roughly determined by the nature of the product (size, shape, etc.). Interestingly, from looking at the current packets (figure 1) Jim has inferred that the customer thinks that photography/pictures are indicative of 'premium' and graphics of 'value', through comparing and contrasting the packets to see what differs amongst them.

It should be becoming clear that design does not happen in a vacuum. The brief circumscribes the space, to a greater or lesser extent: make it look value or premium, try graphics or photos, colours or the impression given by the font may be suggested (e.g. in an example below a 'shouty' font was considered appropriate). Jim guides Annie by suggesting what things to do, how to start off (draw an outline, give prominence to X or Y), how to scope things (make it like product X and Y, not like product Z) and so forth. All this provides scaffolding for the creativity.

Knowing how to start designing for a client clearly has a number of social and methodical components. It involves drawing on customer supplied resources, known about collections (e.g. on-line, in the office) and specifically collected resources (e.g. similar products in the supermarket). The designers pick out important elements to focus on in the design, whether that is a message, an image, a font or a mixture of elements. This often involves some reading between the lines or translation of what the customer has asked for. In the example above this involved divining and translating the customer's notions of quality into more concrete requirements. It involves thinking about where the product will be placed, who it is aimed at, and in many cases for these designers, making comparisons in terms of products it is like and not like. They have fairly standard procedures for beginning – often producing an outline, and making visual searches of font or image libraries. An element of creativity is making the product speak – it should project an image, say something about itself, at the same time being harmonious, attractive and so on (see figure 2 for Annie's redesign).

As well as working within the context a singular product and the customer's vision of it, many products are part of a range and this introduces the constraint of harmony across the range. The designers talk about envisaging ranges of products all sitting together on a supermarket shelf – and looking good together. Given their propensity to do this imaginative work as part of their day to day design work – 'it's for little boys, in value shops, in a lunch box etc.' – it does not seem like they are simply trotting out of a cliché and one of the products they often produce is a 2.xD illustration of a set of product designs. An impressive example of this is that done for a range of Indian food products, curry pastes, flatbreads, spice mixes etc. (see figure 3), set up like a supermarket shelf.

Of course, designers often design to some extent with a notion of extensibility that may be more or less defined and be more or less explicitly oriented to in their work (of course it may manifest itself, sometimes unfortunately, as a backward compatibility that must be achieved – 'oh no now they tell us we need to come up with five more packets and colour schemes to go with the original').



Figure 3: Curry 'range' illustration

As we work through further fieldwork examples the way in which the design context both inspires and constrains the design will be further elaborated.

Specifying the Design: Finding, comparing, choosing

Designs are made up of a variety of elements placed together in a harmonious way to create some form of impression of, or message about, the product. Current technologies enable the designers to do large-scale searches to collect together possible elements of a design. They go through a process of discovery and trial and error, collecting and comparing which leads to a choice of elements from which they can build up their design. A crucial element of this is, unsurprisingly, looking but as we shall see the different technologies enable different ways of

looking. All the various elements that make up a design are important – the colours, the pictures, the fonts – and the designers may search for these individually, but with an eye to how the whole product will look.

In the next example, we will examine how the designers go about choosing the fonts for a chocolate bar redesign (see figure 4) – using websites to select possible fonts and then graphic design tools to compare them. As with the previous example, Jim first instructs Annie, who will do the initial redesigns.

1. J – Just do quite a few Bites he quite likes the colours... as Jane says it quite looks like Lion Bar you know
2. A – Yeah 100% yellow
3. J – Cheap and cheery like what (a) mars bar looks like, what lion bar looks like, it's going to be just like a nice shout-y name for bites.
4. A – Does he want it fun like that cartoon-y type
5. J – Yeah, yeah, yeah
6. A – So he's not looking for that traditional sweet shop
7. J – Nothing traditional, it's going to be going into the pound shops, you buy a wrap, like that, again it's going to be a pound, and the idea is you think – oh my god – I can get all of those chocolate bars for the kids lunchbox or whatever, they're not going to call them lunchbox packs, cos hehe they're not healthy for a lunchbox. Rather than buying a mars bar that costs 37p ... but you know they'll all say bites, they'll say bites, bites, bites, like a packet of breakaway or something like that

Again in this brief, Jim makes comparisons with other related products, he also talks about the colours, and the context of the product presentation (where it will be sold, how it will be presented, what its image is to be). From this discussion of the brief, and the original product, we can see that the product name 'Bites' is going to be a central element of the design. The client wants a 'shouty' name (turn 3), Annie clarifies asking if the font is to be 'fun' and 'cartoony'(4) and when talking about the presentation Jim emphasises 'bites, bites, bites' as the product will be packaged in sets of eight (he demonstrates this here by holding up an 8 pack, but of plain white packaged bars).

Annie starts the design by tracing round the physical packet by hand. She then adds some more lines to the drawing by hand to give it a more 3D effect and to demarcate the panel for the brand name. She then scans this onto her computer and uses this to-scale outline as the 'canvass' for her design (interestingly this mixing of paper and digital work has also been observed with more traditional artists (Eales, 2005)). Next she sets about selecting fonts. To do this she accesses the Dafont website which has a wide range of different fonts, each one illustrated as name of the font spelled using the font. They are a collection of 'amateur' and professionally designed fonts. Not all (especially 'amateur' designed) fonts have a full alphabet of upper and lower case letters. They may just have one alphabet of either, or a mix of the two. They vary in quality. They can often be used for free or for only a very small charge. Tools such as The Internet enable this wide-scale sharing of designs – anyone can create fonts and make them available for others to use – for free or a fee. This gives designers a much wider choice of elements.

Annie searches by looking in various libraries. In the website the fonts are organised according to broad categories (e.g. fantasie, techno, gothique³) and sub genres (fantasie has cartoon, comic, groovy). Within fantasie and in various sub-genres Annie rapidly scans through series of fonts focusing on different examples and for certain ones she opens up their alphabets. Through looking and thinking Annie produces a palette of around 10 fonts. She then spells out 'bites' in a number of the promising fonts on the packet outlines (see figure 5). Selection of the font is a matter of trying out the brand name in a variety of fonts and seeing how it looks. Later on the ethnographer asks Annie about her work:

1. E – So, Annie, when you're choosing like a font, what makes it like the right font? Is it partly to do with, the word, and the packet?
2. A – I think for me when I'm looking at it, I don't know if it's a combination of who the market is, and, how you feel when you look at it really, like, for example, I really like that B, there but everything else being in capitals just ruins it... I think laying them out like this and you can see them all on the packet, allows you to kind of sit back and go actually which ones do I like and you choose from them maybe three
3. J – That that's great example of how fonts look and stuff on, some of the different things you're getting off those is mad, the fact that the one at the top, right, doesn't even look like it says bites even to me, It looks like it says bee-tess or something like that
4. N – It's French for bites



Figure 4: Old 'bites' designs



Figure 5: New 'bites' fonts being tried out

Annie begins (turn 2) in response to the question drawing on a notion of market appropriateness before focusing on 'how you feel when you look at it'. A notion of 'market' drove an initial focus on 'shouty' and 'cartoony' fonts, however, much of this work seems to be about looking and trying out and seeing what you think. On-screen comparison allows the designers to evaluate the fonts. Annie firstly draws attention to the second from top bar on the left, and states how she likes the 'b' but everything else being in 'capitals' causes a problem. In the photograph pictured the 'b' is now also a capital letter, but this is only because it

³ It is a French website – i.e. fantasy, techno, gothic.

has now been ‘paint-shopped’⁴ into one by Annie, since there was no capital ‘b’ in the font alphabet. Hence, with the original there was a fairly obvious lack of harmony between the letters, for this name and product at least. Jim chimes in, in agreement and then draws attention to a problem he sees with another bar, top right – it does not look like it says bites but ‘bee-tess’. For him, the ‘t’ sticking up breaks up the flow of the word. This brings in the notion (commonly discussed by designers) that a good font looks good and visually ‘says’ the word and the brand nicely. Jim then continues (see below) to suggest his preference – second top, right hand side – it is the right kind of font, of that market, and again he draws comparisons, *Boost* and *Sainsburys* homemade. In this turn he also discounts the font below it for a lack of consistency (harmony) in letter size.

5. J – The one below it, is really, of that market, it’s really perfect as something that you’d expect to see, you know like as a Boost or something like that... the one below that, it’s quite a nice font, but it will never work with *Bites* because your *e* and your *s* are massive compared to your what look like a lower case *b i t*... it (back to the one above) looks like a Sainsburys, homemade, chocolate biscuit bar you know doesn’t it really, it’s just got that rounded, sort of regular, it’s quite normal
6. A – It’s chocolaty
7. J – The one that’s, the really fancy one that’s underneath the one you said you liked the *b* on, on the left hand side that’s doing something in another way, that’s saying bites to me, BITES as a like a fierce word
8. A – For boys, yeah, little boys
9. J – It’s quite quirky, it’s quite doing it, yeah, it’s saying the word bites without it just spelling the word bites, its got a bit of a crunch to it ain’t it

Annie agrees with the favourable assessment of right hand second top – making a comparison with the product substance itself, ‘it’s chocolaty’. Jim then goes on to evaluate left side, second from bottom – it does something different, ‘BITES as a fierce word’. Annie suggests this makes it something for ‘little boys’, and Jim continues that it spells bites in a ‘quirky’ way with a ‘crunch to it’.

Just as the designers might be considered to create a palette for the fonts, they also more conventionally create palettes for colours, although they use it in a different way. In this next example, we examine how during a redesign for a range of puddings, Just Puds, the designer, Neil, creates and uses a colour palette. Furthermore, we are given an explanation of how different colours can represent different types of products in design. Applications like Illustrator come with a wide range of custom palettes (which can contain 20-50 harmonious colours and shades) and palettes are also available for view and/or download from the Internet (e.g. see COLOURlovers.com for palettes that may break all the traditional ‘rules’ of harmony). In this case Neil selects a colour palette from the Illustrator library and from it sub selects a set of colours for his palette.

⁴ The Internet allows Annie to rapidly access a large collection of interesting, non standard fonts. Adobe Illustrator© facilitates her alteration of a base element in that it allows close-up fine detailed work in which the letter may be altered in a way that seems seamless. Clearly her artistic abilities – drawing and aesthetic appreciation – also enable her to do this.

1. N - But to start to get there we'd like to evolve. So they like the matt finish and everything it's just lack of, it won't be highly laminated like they are. Errm not important at this stage apart from thinking about how the colours will look without a sheen on them. And the luxury feel of the packaging, they don't want the dark stuff because it's a bit too luxurious possibly, they prefer the brighter colours. So from that step I thought it would be a good idea to choose a colour palette
2. E - Yeah this is what I was looking at you had a palette yeah
3. N - Cos it sounds like they want to do it as a range, they've already got a range, they want to fit in with that (shows me current product range) and take their product
4. E - Is that their current ...?
5. N - That's their current list of products and I've just got sort of some colours that are fairly near but probably a bit richer

The notion of the range is clearly important in the way that designers do their design. As Neil produces the re-designs for Just Puds (see figure 6) it is clear that he has a palette of colours (turn 2 above, and visible in figure 6 as a set of colour swatches close to the designs, which are taken from a custom palette open on the right hand side of the screen). Neil has picked a palette in reference to the current colours of the Pud boxes but 'richer', 'deeper' etc. This was also clearly about correspondence to the flavours of the desserts – chocolate, ginger, caramel, toffee, lemon, plum – so yellows, browns, oranges, reds, purples. Neil is seeking an appropriateness of colours that work in various ways – appropriate to the flavours, generally in keeping with the colour categories across the range for the old designs, appropriate to the customer's wishes of a new direction and, finally and crucially, appropriate to each other (harmonious).

Watching Neil doing this task, it is apparent how he builds up from one box – an assembly of background, text, and picture – to another. With the first design, the selection of colours, font and photograph and the resulting internal harmony between these elements is the primary focus. As he designs further boxes, one-by-one, Neil then evaluates not only the internal harmony of elements but the harmony across the boxes. He re-adjusts colours (and all other aspects of layout), doing new versions and so forth. Work on whatever box is next, invariably becomes work on the whole collection. For example, the top three on the left are 3 versions of the same pudding, produced first then altered twice in response to unfolding constraints produced by subsequent puddings. Harmony across a range produces a series of constraints that are both internal to the individual item and across the products in the range. They are dealt with as a set of local adjustments, and luckily it is not as if there is a single solution, rather, a solution is arrived at through trying, looking and adjusting.

The 'bright' colours of the customers request are translated into 'rich' colours in the design, showing how the brief helps shape the design space but does not dictate it exhaustively. In the extract below, Neil is discussing these colours with the ethnographer and we can see how language is used to articulate the design – to elaborate what the design is 'saying'.

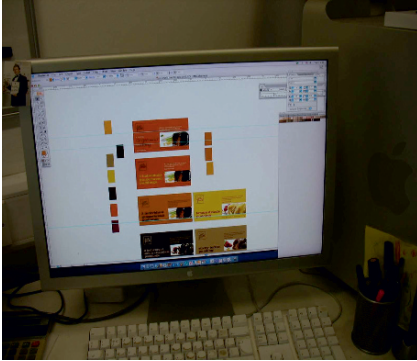


Figure 6: Harmony in the Just Puds range




Figure 7: Highland Cow toffee picture

1. N – Things like that. This is, this is, I should imagine that they're not cheap... and they want it to be seen as a luxury treat... a real luxury indulgence... so the colours can't be that bright, I think what they mean by the bright colours, they probably mean rich as well
2. E – Yeah
3. N – Not just bright, there needs to be a rich brightness to it so if it's a red it's going to be a deep
4. E – Yeah
5. N – Luxurious
6. E – Yeah
7. N – Or a crimson rather than a red
8. E – Yeah
9. N – (showing the red) That's you could go red
10. E – Yeah
11. N – Or you could go red and mean it... you know what I mean... richer erm, which if you go and choose a rich colour you're immediately explaining the product because the products meant to be rich, as well... you probably have to be rich to afford it, as well but... but that's what they're trying to put across, they're trying to say. I mean they're saying here they want it to be seen, as a 'aimed more at an upmarket delicatessen, Harrods, Fresh and Wild, rather than bigger supermarkets, so that's what I'm trying to, really, attain

Neil, as we have seen before, scopes where the product will go – 'Harrods', 'Fresh and Wild' (turn 11) its niche - the 'luxury' end of the market – and, therefore, that it should have 'rich colours'. There is repetition and the choice of particular descriptors (luxurious, rich colours), and the contrasts of what they are not. Neil builds up, turn on turn, a 'vision' of the product in relation to the red colour on screen; 'deep, luxurious, crimson rather than red, you could go red, or you could go red and mean it'. It may be somewhat of a moot point whether this red is inherently and unequivocally deep and luxurious but the rhythmic qualities of the delivery, the emphasis on the three words and the comparison to 'normal, average, red' together serve to encourage the viewer to agree with the description being offered. It certainly has features of a 'sales pitch' in its efforts to convince!

In the example below (continuing the previous) we can see how the designers apply their design knowledge in making creative decisions as Neil defines how certain colours go with certain products.

12. N – Just trying to, because they are on about wanting to be upmarket, they want bright colours but for a bright colour to be upmarket it can't be a garish, cheap looking colour
13. E – Um hum, and you mentioned the matt finish as well
14. N – Yeah
15. E – What would take a matt finish well
16. N – Um we want, the thing is with bright colours you've got to be very careful not to make it look like, em, confectionery, children's... or dairy products, I mean a lot of dairy products are brightly coloured...soft drinks... juices

Neil suggests that there are very particular constraints in finding a bright but upmarket colour – there is a danger of it looking garish and cheap. Bright colours are more often associated with children's (foods), dairy products⁵ and soft drinks. We saw above that bright, in this case, becomes 'rich, luxurious and crimson' meaning a deeper darker red. Indeed, throughout all of the examples discussed so far we can see designers employing a 'grammar' and knowledge of types of food to particular colours – which is fairly closely followed in the case of the Bites Bars and Just Puds. Chocolate is associated with brown, gingerbread with orange, strawberry with red, caramel with yellowy-brown and, slightly more unusually, coconut with blue (although not for those familiar with '*Bounty*'). Of course, while individual products 'suggest' certain appropriate colour categories, finding the right shade of green, for example, is seen as crucial.

Another interesting feature of design is the way in which search proceeds. Commonly search is vague and the act of searching is about looking for inspiration, rather than retrieving something you know is there, although it is guided in some way – in part by the constraints of the design ideas, as the Bites font search was. Another interesting feature of search is that various possibilities are selected and often the designer does not know which picture etc. will be used until further work, talk etc. has been undertaken. In another example, shown above (figure 7), the designer was re-designing boxes for toffees, fudge etc. and following on from the company logo thought that an element of the box should be a highland cow: but in what way? An extensive search of graphics and photographs was undertaken with a number being selected and tried in various combinations with other elements of the design. The final photograph selected is a mixture of two – one of the highland cow; the other of the mountains in the background. Blending the two photographs together such that they look as one requires a lot of detailed skilful work. As well as doing the Photoshop work for aesthetic reasons, putting together and altering photographs from large on-line collections like this allows the designers to create a unique brand specific image much cheaper than it would cost to get a bespoke photo taken by a professional.

⁵ If you look at milk bottles, yoghurt cartons etc. they very commonly have bright basic colours – blues, reds, greens, yellows - certain 'looks' and colour choices are favoured by different types of products.

Constructing the Design: Editing

Looking (cf. Coulter, 1990) is not a single order of activity. It is crucial for appreciation and editing. It is motivated to doing different things and it is conspicuously different in form. The designers regard a design from further away, inspect details closer up, both by looking close up to their screen and enlarging (e.g. a letter) massively. They gain comparison and perspective by looking at designs together, on different formats, in different lights and obviously by sharing them and discussing them. Some of these ‘lookings’ are just on-going *thinking about* (as a perspicuous visual activity, mundanely identified) or assessing the design. Some are more obviously to do particular things like the close up work to turn the ‘b’ into a capital, which requires fine-grained blending using a painting programme at something like a pixel level, but the new capital letter also needs to be looked at further away to judge it from some perspective. There often is a to-ing and fro-ing between different views and different ‘lookings’ – something should look good ‘from all sides’.

Another feature of designer expertise is their appreciation of harmony and their understanding that harmony (amongst constituent parts) can only be ‘accurately’ judged when viewing the assembled product. This is less a theoretical position, than an inherently practical one, and it is clearly demonstrated in the way designers work. Just as we saw with Bites and Just Puds, definitive decisions on many features of the design are delayed. Designers tend to keep some options open. Early in design a number of parallel choices (e.g. of fonts etc.) are investigated and compared. The tools support this – enabling designers to mock up, and change rapidly, products and ranges, seeing how the colours and fonts will look within a product and between products. Designers gradually build up different elements in a layer-like fashion (outline, font, then colour is common, but it depends on the product as graphics or photographs may be crucial elements). During these steps there is some proliferation but also elements and ideas are discarded. It is clear that an element such as a font or colour or photograph that at one point may have been a favourite will be rejected later on because it does not work with other elements, just as a previous rejected or ignored possibility may make the final project. The tools – both the web based libraries and the design technologies themselves such as Paintshop support this ‘massive experimentation’ – enabling designers to skim through a whole range of different elements to use as starting points, and mock-up designs and ranges relatively easily, particularly just through version saving, copy and paste and various specific tools like palettes.

Harmony can only be assessed in an assembled product and achieving harmony in a product often means adjusting the internal elements, or across the range. Using palettes enables some prejudging of this in advance, but does not preclude needing to try out the colours etc. and making judgements based on the actual designs. As the design progresses, the designers can then add in the other

elements, make small alterations or major changes, copy and paste and adjust the elements to suit the next package in the range and so on.

Discussion: Technology, practice, talk and creativity

We have endeavoured to show how 'creativity' can be understood in quite mundane ways as arising out of the mundane technologically-supported collaborative practices that go into creating packaging designs. Nevertheless, we are not trying to suggest that is 'all there is', for the practices we describe above are visibly and accountably 'artful'. The work is, even so, methodical, in that it is carried out according to known and shared patterns of activity, practices and resources. It has a 'loose' iterative workflow – create outline, create palettes, mix and match and so on and it is tightly coupled with (supported by and shaped by) the technologies they use. Search libraries enable rapid, widespread and opportunist discovery of design elements: images, fonts, even colours. Once they have chosen elements they can import them into their design tools and easily experiment with combinations, alterations and so on. Although we have said that the tools support this kind of work well, this does not mean it is not painstaking – designers spend time getting the effect just right, whether it is an individual letter (cf. the B of bites), an image (cf. the highland cow) or an assembly of parts or of the range. The tools help to enable searching, comparing, zooming in and out, editing, creating palettes, discovery, emergent effects, and so on. The designers can try things out, see the whole range, and so on. The projects and the artefacts they create – at all stages, from looking through libraries to almost finished designs – form the basis of their cooperative work. They are available for the other designers to oversee, they become objects for discussion – the comparison is often done cooperatively- they are 'worked up' through talk into being closer or further away from what the customer wants, from what the product should say and so on.

One reading of this elucidation of some of the social and methodical aspects of creative and aesthetic work would be that it has provided an understanding of creativity that opens the 'black box'. That is, in this context we see quite ordinary and routine features presented as part of the creative process. Some of these features are 'scenic', insofar as they constitute a background orientation that informs a large part of the work. In this case, the main scenic feature is the customer, who is present in various ways - in conceptions of the 'brief', the 'market', the 'range', the 'brand' and so on. Others, we can quite clearly identify as a series of easily understandable practical activities – decisions on how to proceed draw on resources and knowledge that once put into a design make logical sense. We see how elements of customer products and requirements are seized upon to scaffold the design and make decisions about which way to proceed. Other products similar and different are brought into play to delineate

and shape the design space. Developing designs are compared and contrasted to see their fit with the requirements. Reasons for deciding to proceed in one way or another, or valuing one design over another are to be found within the process, in the talk of the practitioners, and can be provided readily after-the-fact, as *the* sensible way to proceed. Assembled in some way, the steps we have outlined above are more or less necessary elements to the design process in this context. But they are evidently not sufficient, for if they were it would make no sense to speak of creativity in the first place.

Our point (again) is that these assemblies are *artful*- they involve the deployment of shared knowledge, resources and histories that constitute the practices of assembly, elicitation, specification and construction outlined above. Creativity and aesthetic judgments are not especially mysterious - they are manifested in knowing which type of resources to marshal to scaffold the design, how to pick out of the possibilities they provoke, how to assess the on-going design, understanding why one thing 'works' and another does not and making the 'good' choice. Whatever it is that originally 'sparks' creativity, designers willy-nilly have the everyday knowledge of how to proceed at given points in the design process, what sorts of things to do, what customers 'like this' may be looking for, and so forth. One last element we should make explicit is that the aesthetic skills entailed are evidently 'talkaboutable' (cf. Turner, 1969) - a critical feature of this collaborative design work. 'Talkaboutability' plays a special role in this process, implicated as it is in the work of brainstorming, directing, and choosing. It assists in providing depth to the evaluations: something is 'of the market', or it's 'shouty', or 'quirky' or got a 'crunch to it', or 'for boys', or 'chocolaty', or says 'bee-tess'. In this kind of talk we witness an elaboration of what is seen, an explanation (e.g. of why something 'works' or not), or an invitation for the viewer to see the product as such, in terms of who would eat it, a sound, a texture, a style - in other words some greater context within which to view the object, or that it can be related to. The comparisons are relevant in terms of other elements of the experience of it e.g. as a food product ('sensually' its texture; crunch, chocolaty) for a potential audience, referencing a style (quirky, shouty). This talk performs a vital role in sorting out what a design 'says', deciding what designs are most liked and reaching an articulated agreement. It is important to recognise that the articulation performs an important role in terms of sharing a preference, of emphasising it. The collaborative 'looking, articulation and discussing' element is very important in sharing an idea or an understanding, and it also serves as an on-going assessment of ideas. To return to Wittgenstein (as quoted at the start) graphic design is a 'form of life' within which aesthetic talk is abundant. The specific meaning of the words chosen is not inherent and objective, without context, but is to be located in their 'gestural roles in complex sets of activities'. They are understandable as purposeful 'moves' in the activities of creating and evaluating design. Furthermore, the activity relates heavily to later 'presentation

activities' – talking about designs is a key skill of graphic designers when they present to clients. In those situations the convincing description of what the client is seeing, what the design suggests and why it fits the company-product is very important to the 'pitch'. This discourse is built up and refined throughout the design process, such that it provides a strong basis for the polished presentation to the client at the end of the process iteration. The narrative of what a product is about, who it speaks to, what it says etc. is built up in the trajectory of the design process, often through several iterations that tend to lead towards a refinement.

There is, as stated earlier, some kind of 'grammar' of design (what things customarily mean and how they go together) and 'stock' resources of design knowledge. Designers draw on both in making design choices about what kind of font, colour, overall 'look and style' would fit for a certain type of product aimed at a certain type of market. Their discourse is punctuated with descriptions of what is 'expected', 'bang on' or 'reasonable' in terms of the elements of the products they are designing. According to theories of design these might be called the 'cultural referents' of those elements and those designs (i.e. they visually 'name check' other things in a culture that contain those elements). Another interesting feature of this is the notion of fashion, and its complexity and self-referential character. Hence, whether a product is meant to be 'retro' or 'classic', or 'premium' or 'innovative and ground-breaking' or 'value' will have profound implications on how the grammar and knowledge will be drawn upon.

In sum, this paper has been concerned with mundane aspects of creativity. It sees the creative process in graphic design as being mutually constituted in orientation to certain 'scenic' features which pertain to customers and markets; in certain kinds of visible 'orderings' which are done in artful ways using technological resources, and in the talk which constructs the 'cultural referents' visible in their work. Lest it be thought that these matters are entirely local to graphic design, we should perhaps remind ourselves that art history and the sociology of art produce similar (if occasional) reminders about the fine arts. Clement Greenburg's (1992) work on modernism was very much about how language (written text and talk) about art transforms the art. Howard Becker (1982) showed very persuasively how art production was bound up in a variety of collaborative market practices. We would anticipate that some features of graphic design are common to other areas of creativity. Some will be quite different. If problems of similarity and difference characterise much of the conceptual work done in CSCW, however, then hopefully this paper provides an initial springboard for the analysis of the visible aspects of the 'creative'

Bibliography

Ackermann, M., Pipek, V. and Wulf, V. (eds.) (2003). *Sharing Expertise: Beyond Knowledge Management*. Cambridge, Mass. MIT Press

- Baxandall, M. (1974). *Painting and Experience in Fifteenth-Century Italy: A Primer in the Social History of Pictorial Style*. Oxford Paperbacks, Oxford
- Becker, H. (1982). *Art Worlds*. Berkeley: University of California Press
- Benjamin, W. (1982). The work of art in the age of mechanical reproduction in *Modern Art and Modernism*, eds. F. Frascina and C. Harrison, Open University Press: 217-220.
- Button, G. and Sharrock, W. 1997. The Production of Order and the Order of Production, in J. Hughes et al (eds.), *Proceedings of ECSCW '97*. Dordrecht. Kluwer
- Cook, E, Teasley, S.D. and Olson, J. (2005) Heterogeneity in Harmony: Diverse Practice in a Multimedia Arts Collective. GROUP, 2005. 334-335
- Coulter, J. (1990). The praxeology of visual perception. *Inquiry*, 33: 251-72
- Dourish, P. (2006). Implications for Design. Proceedings of CHI 06. Montreal, Canada. © ACM
- Eales, R.T. (2005). Creativity in action: some implications for the design of creativity support systems. In proceedings of CHINZ 05. Auckland, New Zealand
- Fischer, E., (1970). *The Necessity of Art: A Marxist Approach*. Harmondsworth. Penguin
- Fisher, G. (2005) Distances and Diversity: Sources for Social Creativity. *Creativity & Cognition* 2005. 128-136
- Goodwin, C. (1994). Professional Vision. *American Anthropologist*, vol 96: 606-633
- Greenberg, C. (1992). *Art and Culture: Critical Essays*. Beacon Press
- Greene, R. T., (2001). A garbage can model of creativity – the four cycle model – derived from a model of 42 models of creativity, *Journal of Policy Studies*, No. 11. September: 1-204.
- Hughes, J., King, V., Rodden, T., Andersen, H. (1994). Moving Out from the Control Room: Ethnography in System Design. Proceedings of ACM CSCW'94, pp. 429-439
- Herrman, T. (forthcoming 2009). Support of collaborative creativity for Co-located Meetings, in D. Randall and P. Salembier (Eds.), *From CSCW to Web 2.0*, Springer.
- Jacucci, G. and Wagner, I. (2007). Performative roles of materiality for collective creativity. *Proceedings of C&C'07*, June 13–15, 2007. Washington, DC, USA: 73-83
- Randall, D., Harper, R. & Rouncefield, M. (2007). *Fieldwork for Design: Theory and Practice*. Springer Verlag, New York.
- Sacks, H. (1995) *Lectures on Conversation*. Blackwell.
- Schmidt, K. (2002) The Problem with 'Awareness': Introductory Remarks on 'Awareness In CSCW', *Journal of Computer Supported Cooperative Work*, Vol 11, Nos 3-4: 285-298
- Schneiderman, B. (2000) Creating creativity: user interfaces for supporting innovation. *ACM Transactions on Computer-Human Interaction (TOCHI)*. Volume 7, Issue 1. March
- Turner, R. (1969). *Occupational Routines: Some Demand Characteristics of Police Work*, Unpublished paper presented to the Canadian Sociological and Anthropological Association Conference, Toronto, June, 1969
- Whalen, J. (1995) Expert systems versus systems for experts: Computer-aided dispatch as a support system in real-world environments, in P. J. Thomas, *The Social and Interactional Dimensions of Human-Computer Interfaces*, Cambridge, Cambridge University Press
- Wittgenstein, L. (1966), *Lectures and Conversations on Aesthetics, Psychology, and Religious Belief*, ed. Cyril Barrett. Oxford: Basil Blackwell
- www.COLOURlovers.com
- www.dafont.com
- www.delicious.com
- www.flickr.com
- www.gettyimages.com

Using Annotations in a Collective and Face-to-Face Design Situation

Sylvie Guibert¹, Françoise Darses² & Jean-François Boujut³

¹ Equipe d'Ergonomie, CRTD, Cnam, Paris

² CPU Group (*Cognition, Perception, Usage*), LIMSI, CNRS, Orsay

³ G-SCOP, INPG, Grenoble

¹sylvieguibert@orange.fr, ²francoise.darses@limsi.fr, ³jean-francois.boujut@inpg.fr

Abstract. Allowing a group of users to produce and transmit some annotations in common digital documents is nowadays a major issue for groupware systems. In this paper, we report a psychological and ergonomic study carried out on this topic in the mechanical design domain. We observed a collective design process that took place in a series of face-to-face meetings attended by the members of a design team. Our results show the minor role played by textual annotations, contrasting with the great number of figurative annotations. We also highlight that the function of annotations is not to develop parts of the solution but to provide the team members with contextual descriptions of the problem and the solution. These results are a first step towards a model of annotations in a collective face-to-face situation. They also provide interesting tracks for elaborating specifications of annotations in mediated situations.

Introduction

One of the current research objectives in CSCW and engineering design (Computer Assisted Design - CAD systems) is to develop digital annotation functionalities for groupware tools (Cadiz, Gupta & Grudin, 2000; Koivunen & Swick, 2001; Boujut, 2003). The aim is to allow the users of such systems to draw up, transmit and use annotations that are traditionally produced on paper and

created spontaneously during individual or groupwork. Most of the research concerns the field of digital documents, the idea being to use the properties of annotations to make web-based information retrieval more efficient. Studies are also being carried out in the field of CSCWriting (Cerratto & Rodriguez, 2002 ; Weng & Gennari, 2004) in order to provide tools to assist the collective drafting of texts. As regards the field of engineering design and architectural design, the use of annotations is a very frequent process in non-mediated activities. On the other hand, digital annotations provided by certain CAD systems, remain little used, as there are not well suited to the action goal of the designers.

Whatever the field concerned, it can be seen that most of these studies have adopted a psychological and ergonomic analysis of the cognitive process of annotation. Our study aims to contribute to such an approach. On the basis of a state of the art, we have built a framework for defining annotations. Using this framework, we carry out a psychological and ergonomic study in the mechanical design domain. The design process that we observed took place in a series of face-to-face meetings attended by all the members of a designers' team. The results that we obtained take into account the role that annotations play when made in such meetings, and indicate some interesting functions that mediated annotations could meet.

The Use of Annotations in Collective Elaboration of Documents

Computerization has modified the status of a document: whereas this status was previously acquired when the drafting and editing of a text had been completed, the status of "document" is nowadays attributed to texts that are still being drawn up, and whose degree of completion may vary greatly from one situation to another. Annotations that are mediated and shared play a vital part in the progressive breaking down of a hard and fast division between the "completed document" and "document in progress", as they are one of the vectors for the continuing evolution of texts.

Many research studies (Denoue, 2000; Wolfe, 2002; Jacobs Reimer, Brimhall, Cao, & O'Reilly, 2009) aim to improve the use of annotations made on texts in the framework of individual activities by one (or several) readers: processing the information contained in the document, obtaining an automatic summary and information retrieval being the most widely examined functionalities. Yet annotations are increasingly being used in group situations as, due to the digitalization of annotations and their widespread diffusion, the document becomes both the object of the collective work as well as the collective workspace itself. By examining the few studies that have analyzed real-life situations where annotations

are produced and used, we have identified several objectives that are pursued in this direction.

Reducing the cost of reading a document – A text can be read more quickly thanks to annotations made by another reader (Marshall, Price, Golovchinsky, & Schilit, 1999). This presupposes that both readers have the same reasons and context for using the document.

Reducing the individual workload for processing the information – The members of a group can pool the resultants of their documentary research. Making the information jointly assimilable is done by indexing the documents and will increase the field of knowledge open to all (Koivunen & Swick, *op.cit.*; Denoue, 2000). In this case, and in contrast to the point raised above, the benefit is an increase in information for each person involved.

Assisting decision-making – It has been stated (Marshall, Price, Golovchinsky, & Schilit, *op.cit.*) that annotated passages are the most discussed collectively. Thanks to annotation zones (underlining for instance), it would then be possible to identify areas of consensus or divergence on the importance of parts of the text.

Assisting the synthesis of documents read by different persons – It has been observed by Denoue (2000) that the various passages that are annotated by several persons could form a sound basis for producing a synthesis or a summary.

Increasing the relevance of the results of information retrieval – Combining the annotations made by different persons makes it possible to index the document on the basis of unanticipated objectives for using the document (Champin, Prié, & Mille, 2000).

Cognitive and operative synchronisation of the members of a work group – Annotations are used as a means of coordination, as stated in Denoue (*op.cit.* and Marshall, Price, Golovchinsky & Schilit (*op.cit.*): establishing what the members have already read, allocating tasks, planning goals, etc.

Nonetheless, these mediated practises of collective design, based on the sharing of annotations, have a certain number of limitations:

- The annotations that are exchanged are more helpful for coordinating tasks than for joint and direct cooperation of a common artifact, such as text, mechanical artifact, etc. (Cerratto & Rodriguez, *op.cit.*; Churchill, Trevor, Bly, Nelson, & Cubranic, 2000).
- Personal annotations – made by an author for himself – are generally confused with annotations that are intended for the group (Cerratto & Rodriguez, *op.cit.*). This often makes the exploitation of annotations too costly.
- The widespread diffusion of annotations partly diminishes their social, informal and sometimes confidential nature. Some authors consequently limit their annotations and thereby impoverish the content of the exchanges, as stated in Cadiz, Gupta & Grudin (2000).

The Use of Annotations in the Particular Case of Product Design

The documents used in product design are mainly three-dimensional digital representations and geometric in nature. In contrast, the constraints of the work are mainly expressed in a non-geometric way (the nature of a material, the manufacturing process, etc.). The digital transfer of these annotations is not currently possible, as available CAD tools only allow for annotations as simple pointers. One can point out two major properties of the annotations produced and exchanged in the field of design, which must not be impaired by becoming digital and mediated: annotations are “boundary objects” and they support the cognitive synchronization.

Annotations as “boundary objects” — Annotations provide intermediary representations that serve as common frames of reference to support communication between the various actors (and professions) involved. This role is due to the fact that they do not function as closed objects (which cannot be acted upon as, for example, plans) but rather as open objects that are not completely prescriptive. They play the role of “boundary objects” (Star, 1988) - or cooperation entities (Boujut & Laureillard, 2002). Their function is to present the various points of views, specific to each profession and each background, and to provide the members with the means to take part in and support discussions concerning these differences in such a way that a shared understanding may be achieved (Boujut & Laureillard, *op.cit.*).

Annotations as mediums for mediated cognitive synchronisation — Some phases of collective design are not conducive to being mediated, mainly due to the subsequent impairment of the possibilities of oral explanations about certain solution elements on which the designers are focusing. In Blanco & Gardoni (2001), it has been suggested using annotations to make up for the loss of such verbal exchanges. This would, in particular, allow the entire code of the technical drawings to be transmitted. The mediated manipulation of three-dimensionnal objects should thus become operative, allowance being made for the constraints of the design task.

Defining a Framework for Annotations

Although there is not a universally accepted definition of annotations, it is nonetheless possible to establish a number of their characteristic elements, brought to light from the previous state of the art as well as by Azouaou, Desmoulins, & Mille (2003) as well as in Bringay, Barry & Charley (2004). These characteristic elements are summarized in the following figure and discussed below.

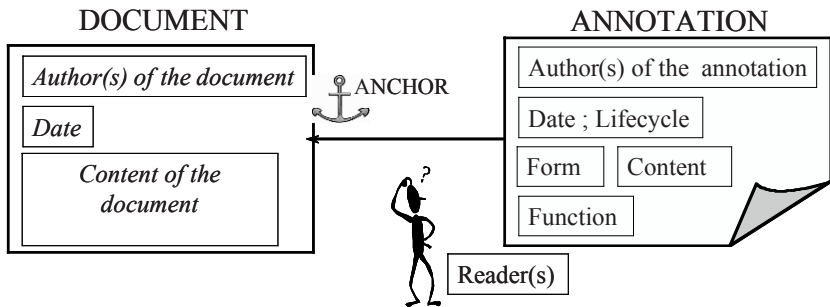


Fig. 1. Characteristic elements of an annotation: a synthesis drawn up from the literature.

- *The document is the target* to which annotation refers. This may concern the document as a whole or only a part of it. The annotation may or may not be physically distinct from the document it refers to. For example, a footnote is integrated in the document whereas a “post-it” sticker is distinct from the document it refers to.
- *The content of the annotation* corresponds to the information that is transmitted. It may be transmitted in various forms (text, icon, graphics, etc.), which are more or less shared by the annotators.
- *The anchor of the annotation* is the informational point to which the annotation is attached. In the example of a “post-it” sticker, the anchor is the place where the “post-it” is stuck, the annotation being the information written on it.
- *Annotations are private or public*, depending on whether or not they are produced in a shared work document and/or made available to a group, or whether they are intended for personal use.
- *The life-span of an annotation* may be of varying lengths: short-term as in the case of a “post-it” sticker saying where to file a document, or permanent such as the footnotes of a document.
- *The annotator and the reader of the annotation may or may not be the author of the document*. For example, a footnote may be produced by the author himself or by his translator. When an annotation is produced, its writer may or may not know the persons to whom he will transmit the message. A teacher who annotates homework knows his student but an editor drawing up a commentary of a book, on an Internet site, is not intending this annotation for any particular person.

Objectives of the study

The functions of annotations in collective design processes are variously considered in current studies. They are characterized more by an intuitive perception of the needs than by a systematic and thorough ergonomic approach. Our study aims to examine this issue in greater detail.

Data relating to annotations were gathered during a collaborative design situation for a product intended for the general public (a bicycle trailer). Four mechanical engineers had to define a manufacturing solution of the product in face-to-face meetings. The meetings were filmed and all documents were collected in order to analyze the annotations produced during the collaborative sessions.

Method

Design Task

The project involved designing a bicycle trailer to carry children (see Fig. 2a). The designers were required to produce the manufacturing and marketable solution of this artefact. Their task was therefore to work on the CAD modelling of the trailer parts, the complete plans and nomenclature of the product, the costs and supplier estimates, together with an implementation strategy. The designers were provided with a file containing the findings of a preliminary study which included a functional analysis of the product, a proposal for a technical solution and some detailed plans of partial solutions.

The design process took place over four weeks. All four designers attended a two-hours meeting each week. Between two meetings, a week was set aside for individual work, during which the designers could only communicate with each other using e-mail. A preparatory meeting was held so that the designers could become familiar with their allocated role and the preliminary study that had already been carried out. After this meeting, and to ensure that the designers could use the CAD tool ProEngineer effectively and had analyzed the architecture proposed, each of the designers had to model parts of the product before the first work meeting. Finally, a review of the project was planned for a week after the last work meeting so that the team could present their work to the client.

Designers' Profile

The project team was made up of four recently qualified mechanical engineers who knew one another but had never worked together. They had experience designing in a workplace setting, and were paid for their participation to this specific design project. They were allocated clearly defined roles: a project leader-

coordinator, an ergonomist-designer, a designer responsible for the connecting parts, a chassis designer. Before the first meeting, the designers received a job description of their roles, together with details of the aims of the project. The project leader was in charge of ensuring the consistency of the solution and the work schedule. The ergonomist-designer was responsible for the products' usability. The two other designers were responsible for the technical and industrial aspects of two separate parts of the product: the chassis and the connecting parts.

Work Space and Data Gathering

The design sessions took place in a meeting room (see Fig. 2b). The tools available to the designers were a Teamboard (an interactive whiteboard), a CAD tool (ProEngineer), a computer with a set of office software packages (Word, PowerPoint, etc.), an Internet connection, a scanner, and blank sheets of paper (A3) and felt-tip pens. The designers had a common storage zone for computer data. The various screen pages of the software packages were displayed on the Teamboard. The designers could interact with them, either by using the mouse or by using a touch screen. The Teamboard could also be used as a whiteboard on which the designers could draw sketches. The designers could draw or make annotations on the displayed pages.

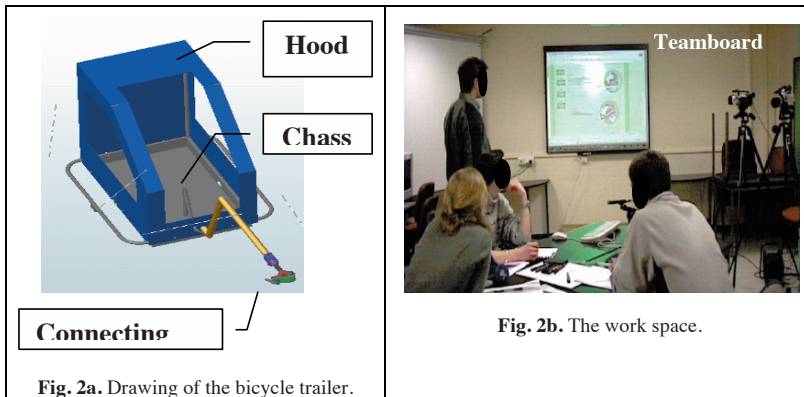


Fig. 2a. Drawing of the bicycle trailer.

Fig. 2b. The work space.

Fig. 2. Design situation

The meetings were filmed with both wide-angle views and close-up shots of certain documents not only to have a global view of the collaborative design process but also to keep a record of the annotations at the time they were produced and to see how they were used. All the documents produced and used were gathered.

Data analysis

Differentiating Document / Sketch / Annotation

We have stated above that an annotation refers to a document. But while it is easy to identify an annotation made on a finished document, it is not so simple to locate an annotation on a document in progress (i.e. a document that is being drawn up). The confusion may, in particular, arise from the fact that certain documents under development, such as sketches, resemble annotations. Therefore, in order to categorize rigorously the various graphical and textual productions and to ensure a sound interpretation of the annotation processes, it was necessary to make a clear distinction between the various concepts prior to encoding the data.

Document. Following Chabin (1997) and Stern (1997), we define a document by its content (the information contained), its container (the medium) and the intentionality that led to its creation (the goal of the activity, the context of the task).

Sketch. This graphical production is a particular form taken at a given moment by a document (or part of a document). It has a context, a container and an intentionality (see Fig. 3) and it may be understood independently of any other document. Any subsequent modification made to a sketch but with the same intentionality is considered to be part of the development of the same document. These modifications (the adding of a line, the changing of a curve) are therefore not considered to be separate annotations.




Fig. 3. An example of a sketch: a content, a container and an intentionality. The medium of this sketch is a blank sheet on the Teamboard (container). It represents a content: the connecting part [1] that joins the bars [2] to the chassis. The sketch was made by the teamleader to explain how the bars can be folded (intentionality).

Annotation. An annotation does not have an autonomous existence independently of the main document. It only reveals its informational content when linked to the document. Annotations rely on a different intentionality to that which led to the creation of the document, which is the medium of the annotation (see Fig. 4).

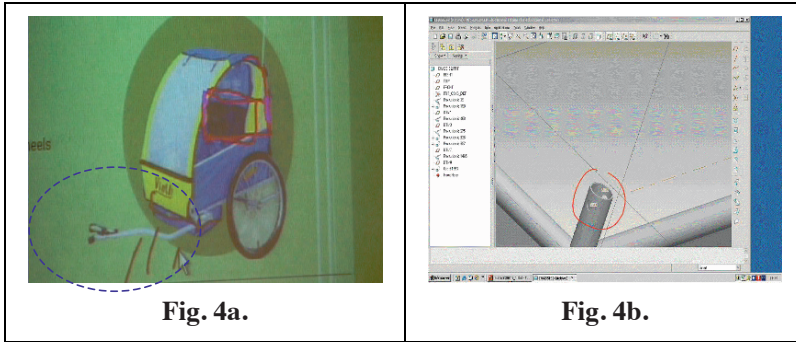


Fig. 4. Examples of annotations, which are both made here on digital screen pages. On Fig. 4a, the annotation is the drawing of a throwing stone under the bicycle. On Fig. 4b, the annotation is the circle which focuses the attention on a sub-part of the chassis.

Analyzing the Data

We produced a chronological representation of the thematic development of each meeting. In order to characterize the medium of each annotation, we also record the way in which the designers used the documents and the various tools according to:

- *the location of the document used* :Teamboard or table;
- *the medium of the document*: type of application for electronic document (CAD, Internet, Word or Teamboard) or type of medium for paper documents (A3 sheets, print outs from Word or journals);
- *the documents' scope of use*: private use or public use.

In this way we obtained a corpus of 93 annotations created during the four design meetings. We also made transcripts of the verbal exchanges associated to the production of each annotation.

Results

Medium of Annotations: A Majority of Digital Mediums

The study revealed that the designers used no private documents, excepting their notebook, which is very often open for all on the table. Almost all documents take place in the collective arena.

It is worth to note (see table 1 below) that electronic documents (48,2%) slightly less than paper documents (51,8%). But two thirds of the annotations (63,4%) were made on electronic documents compared with a third (36,6%) on paper documents (see Table 2). In the study, digital mediums by no means hindered the designers from making annotations, but appeared, rather, to encourage them. This positive

effect of the digital environment on the production of annotations may be explained in two ways:

- *Computer tools are suited to “natural” production:* the digital annotations were produced using a touch pen (supplied with the Teamboard) which reproduces the same functions as a graphite pencil for making annotations: easy to draw with, fast to use. None of the annotations was made using the annotation functions provided by the CAD software, which are not well-suited to such easy manipulation.
- *Increasing use of digital documents downloaded via the Internet:* 32,3% of the annotations was made on documents that had been downloaded via the Internet. These documents concern solutions that are currently available on the market. This approach to design, based on reuse, comes as no surprise, reflecting, as it does, one of the main strategies used in the cognitive processes of design problem-solving highlighted by work on Case Based Reasoning. The hypothesis may be made that easy access to digital documents (via the Internet or Intranet) will encourage and increase the use of these reuse strategies.

Type de support	Documents % (number)	Annotations % (number)	Ratio annot/doc
Paper Documents			
Technical requirements	4,5 (5)	3,2 (3)	0,6
Product booklet published by competitors	6,3 (7)	0	0
Benchmarking documents	3,6 (4)	0	0
Sketches on paper	30,4 (34)	20,4 (19)	0,56
Sketches on personal notebooks	7,1 (8)	12,9 (12)	1,5
Electronic documents			
Digital representations (CAD)	12,5 (14)	25,8 (24)	1,71
Downloaded via the Internet (competing solutions)	20,5 (23)	32,3 (30)	1,3
Digital sketches on whiteboard	14,3 (16)	5,4 (5)	0,31
Technical requirements	0,9 (1)	0	0
	100 (112)	100 (93)	

Table 1. Production of annotations according to the type of the document

It is also worth noting that the solutions produced in the form of digital representations (CAD), although few in number (12,5% of the documents) gave rise to 25,8% of all the annotations, with a ratio of 1,71 annotation per document, whereas the solutions produced in the form of sketches (either on paper or digital sketches) have an annotation /document ratio of 0,56 on the one hand and 0,31 on the other hand.

We interpret this result as being due to the different natures of the solutions in either case. A digital representation describes a solution at a fairly complete level of detail. A consensus among the group has already been reached regarding the solution in question, even if it may be modified subsequently. Any modifications will be marked (pointed out) by annotations. On the other hand, sketches (on paper or digital) represent ideas for solutions, about which the group has made no decision. A modification of such solutions is more easily made by drawing a new sketch than by marking critical points (by means of annotations) on the existing sketch.

The Form of the Annotations: Beyond Simple Deictics

In table 2 below, it is stressed that, of the 93 annotations that the designers produced during the four meetings, only 3 (3,2%) were textual in form (measures, question marks), while 6 (6,5%) were mixed. All the other annotations (90,3%) were graphical. The high number of graphical annotations shown here is not in line with certain proposals for tools to assist mediated design, which are primarily based on textual annotations (in the form of comments). Such is the principle adopted by Naveiro, Brézillon & Soares (2002) in the SISPRO system which enables designers working on civil engineering projects to exchange texts, via a chat room, of all their annotations concerning building plans and design constraints. Our results suggest that such a text-based approach should be enhanced by including graphical annotation facilities.

The graphical annotations produced during the design situation that we studied are presented either as figurative forms (55,9%) or else as deictic (34,7%) – see Table 3. The figurative annotations represent:

- *representation of mechanical parts* (for instance, the position of the seat on the chassis will be roughly annotated on the digital representation of the chassis);
- *pictural representation* expressing problem constraints (the position of the passenger, the throwing up of a stone by the bicycle wheel, etc.)
- *arrows*, which are figuring various elements (the movement of an object, the behaviour of the chassis, the space between the chassis and the ground, distances that must be respected, etc.).

The existence of two types of graphical annotations (figurative and deictic) is an interesting result. Only deictic annotations are available in currently available CAD tools. Our results, on the other hand, show the importance of figurative annotations in the design process. We make the assumption that the rich semantic content of this category plays an important role in the project memory and in the problem solving process, whereas deictic annotations can only have a short-term and shallow role. This being the case, deictic annotations should be completed by

libraries of figurative annotations. The deictic annotations – which account for 1/4 of the annotations – are mostly pointers that serve to highlight certain elements of the problem or of the solution.

Type of annotation	Form of annotation	Semiotic function of the annotation	% (number)
Textual	Text / numbers only	<i>complementing</i>	3,2 (3)
Mixed	Combination of textual and graphical		6,5 (6)
Graphical deictic	Crossing out	<i>erasing / correcting</i>	4,3 (4)
	Combination (Crossing out + component delimitation)		2,2 (2)
	Highlighter	<i>highlighting</i>	18,3 (17)
	Cercle, point, cross, stroke, arrow	<i>pointing</i>	9,9 (9)
Graphical figurative	Movement	<i>simulating</i>	2,2 (2)
	Distance		4,3 (4)
	Component	<i>figuring</i>	46,2 (43)
	Scene		3,2 (3)
			100,00 (93)

Table 2. Distribution of the annotations according their forms.

It is also interesting to note the large number (46,2%) of annotations that represents mechanical parts. Our encoding has clearly established that these annotations are distinct from elements of the solution: they do not represent a development of the solution, but they direct the focus of attention towards a subproblem to be dealt with. We therefore make the assumption that this category of annotations serves to represent the subproblem to be dealt with in a broader context.

The Collective Function of Annotations in Problem-Solving

All of the annotations produced during the four design meetings were created in the workspace, shared and produced by the group. They were public in nature, almost all made on group documents (92%) rather than on private documents (8%), the latter being made available in the shared workspace as soon as an annotation was made on it.

Since most of the annotations are created in and for the group, their function must be understood regarding the collective problem-solving process. To do so, we encoded the annotations in order to interpret them from the collective problem-solving viewpoint. We analyzed the annotations in the problem-solving context on the basis of the verbal exchanges that accompanied the production of the annotation. These exchanges were transcribed, and a chronological depiction of their topic was done (see figure 5 below). The problem-solving nature of the

exchange was also characterized, either as aimed at *generating* a solution or *informing* partners about contextual data related to the problem to be solved.

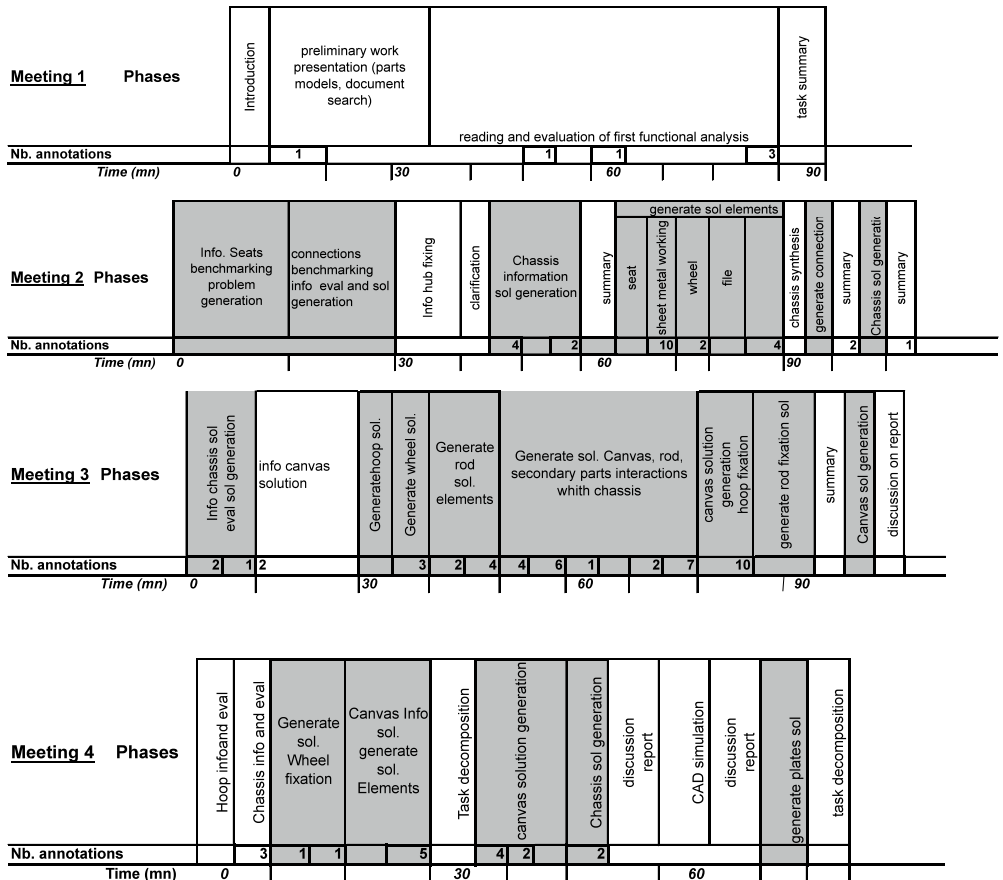


Fig. 5. Distribution of the annotations according to the problem-solving phase during which they are expressed: *Generation* problem-solving phases are figured in shadowed segments. *Information* problem-solving phases are figured in white segments

It was then indicated when – and how many - annotations were formulated along these various phases of the problem-solving process. Results are presented in table 3. It can be seen that most of the annotations (85%) are produced in order to support the process of *generating* a solution. Only 15% of annotations are produced to *inform* the team members about the problem or solution data.

This means that annotations are produced in order to elaborate the solution itself, rather than for establishing a shared context about the problem to be solved and its data. This is an exciting result to be investigated further, since it contradicts some previous studies (Salember & Zacklad, 2007; Naveiro *et al.*, *op.cit.*) claiming that

annotation is mostly dedicated to an evaluation function in a collective problem-solving process.

	Function in the collective problem-solving process	To inform about a solution	To generate a component
Meeting #1		6	0
Meeting #2		3	22
Meeting #3		2	42
Meeting #4		3	15
TOTAL	93 (100%)	14 (15%)	79 (85%)

Table 3. Distribution of the annotations according of their problem-solving function

Discussion and Conclusion

The results presented here are remarkable in several ways. First, they represent the first psychological and ergonomic study that, to our knowledge, has been carried out on the use of annotations in a collective product design situation in which modern digital tools were used. An initial finding is the major role that electronic documents on the Internet play in the problem solving process, compared to paper documents. We put forward the assumption that access to electronic documents encourages analogical and case-based reasoning. We have shown that the production of annotations is stimulated by this availability of analogical solutions. This finding, if it were backed up by studies made in other collective design situations, would confirm the need to provide designers with powerful annotation functionalities.

The analysis that we made of the various forms taken by annotations provides particularly interesting perspectives for the future. Firstly, the minor role played by textual annotations, as observed in our study, undermines the importance that is generally placed on this category of annotation. Textual annotations are supposed to make up for the lack of verbal exchanges that occurs when an activity is mediated. Our results suggest, at least as far as mediated design activities are concerned, that this principle is not sufficient to support the complexity of the activity. Textual annotations will have to be completed by graphical annotations whose semiology is better suited to the work of designers and which provides a less costly manipulation. Secondly, we have brought to light the wide variety of graphical annotations. Although they were traditionally considered as pointers, our study shows the crucial part that figurative annotations play in describing constraints of the problem and the solution. As we have shown, the function of annotations is not to develop parts of the solution, but to provide the team members with contextual descriptions of the problem and the solution, by specifying which data and constraints should be taken into account. Therefore, we do not recommend trying to integrate, in CAD systems, annotations in the digital model of the artifact being

designed: the annotation must be distinct from this model. Its role is to improve the designers' understanding regarding the artifact by specifying the representation of the problem and the solution.

Our study should be taken further in order to examine the possible effect of the function of the designers and the social organization of the group on the production of annotations. It would appear that the more questions regarding the designers' workload are discussed, the more annotations are produced. Testing this hypothesis would require complementary methodological developments. If it were verified, it would be possible to assert that the designers' roles would be an important piece of information to include in the widening of annotations. This possible link between the production of an annotation and the function of its producer must surely be influenced by the social organization of the design project, in which roles sometimes emerge that do not conform to the functions that were attributed at the outset of the project. Nevertheless, such a result would confirm the importance of a multi-viewpoint indexing of annotations.

Acknowledgments

This study was supported by two CNRS programmes: *TCAN* and *Sciences de la Société*. The authors would like to thank the ECSCW reviewers for their stimulating comments on a previous version of this paper.

References

- Azouaou, F., Desmoulin, C., Mille, D. (2003). Formalisme pour une mémoire de formation à base d'annotations : articuler sémantique implicite et explicite. *Actes du Congrès EIAH 2003*. Strasbourg, France, 15-17 avril, pp. 43-54
- Blanco, E., Gardoni, M. (2001). Supporting non structured graphical information in integrated design team. *Proceedings of the 13th International Conference on Engineering Design, ICED 01*. Glasgow, UK, 21-23 August, pp.11-19
- Boujut, J. F., Laureillard, P. (2002) A co-operation framework for product-process integration in engineering design. *Design Studies*, 23 (5), 497-513
- Boujut, J. F. (2003) User-defined annotations: artefacts for co-ordination and shared understanding in design teams. *Engineering Design*, 14 (4), 409-419.
- Bringay, S., Barry, C., Charley, J. (2004) Annotations: A new type of document in the Electronic Health Record. Presented at the *Second International Conference on Document Research and Development in Sciences, arts and business: DOCAM 2004*, University of California (Berkeley, USA), October
- Cadiz, J. J., Gupta, A., Grudin, J. (2000) Using Web Annotations for Asynchronous Collaboration Around Documents. *Proceedings of CSCW2000*. Philadelphia, Pennsylvania, USA, 2-6 December, pp 309-318

- Cerratto, T., Rodriguez, H. (2002) Studies of Computer Supported Collaborative Writing Implications for System Design. In M. Blay-Fornarino, A. M. Pinna-Dery, K. Schmidt & P. Zaraté (Eds.), *Cooperative Systems Design. A Challenge of the Mobility Age*. Amsterdam : IOS Press, pp.139-154
- Chabin, M. A. (1997) La cinquième dimension de l'archive numérique. *Document numérique, Vol. 1, 2*, 205-216
- Champin, P. A., Prié, Y., Mille, A. (2000) Annotating with uses: a promising way to the Semantic Web. *Proceedings of Workshop on Knowledge Markup and Semantic Annotation*. Victoria, Canada, 21 October, pp.79-86
- Churchill, E. F., Trevor, J., Bly, S., Nelson, L., Cubranic, D. (2000) Anchored Conversations: Chatting in the Context of a Document. *Proceedings of CHI 2000, The Future is here*. The Hague, Netherland, 1-6 April, pp. 454-461
- Denoue, L. (2000) *De la création à la capitalisation des annotations dans un espace personnel d'informations*. Thèse de Doctorat d'Informatique, Université de Savoie, France
- Jacobs Reimer, Y., Brimhall, E. , Cao, C., & O'Reilly, K. (2009) Empirical user studies inform the design of an e-notetaking and information assimilation system for students in higher education. *Computers & Education, Volume 52, Issue 4*, 893-913
- Koivunen, M. R., Swick, R. (2001) Metadata Based Annotation Infrastructure offers Flexibility and Extensibility for Collaborative Applications and Beyond. *Proceedings of K-CAP 2001, Workshop on Knowledge Markup and Semantic Annotation*. Victoria, Canada, 21 October
- Marshall, C., Price, M. N., Golovchinsky, G., Schilit, B. N. (1999) Collaborating over portable reading appliances. *Personal and Ubiquitous Computing, 3 (1)*, 43-53
- Naveiro, R. M., Brézillon, P., Soares, R. (2002) Knowledge and cooperation in design : the SISPRO project. *Document Numérique, 5 (3-4)*, 115-134
- Salembier & M. Zacklad (Eds.), *Annotations dans les documents pour l'action*, Paris, France: Hermès.
- Star, S. L. (1988) The structure of ill-structured solutions: heterogeneous problem-solving, boundary objects and distributed artificial intelligence. In M. Huhns & L. Gasser (Eds.), *Distributed Artificial Intelligence, Vol. 3*. Los Altos, CA: Morgan Kaufman, pp. 37-54
- Stern, Y. (1997) Les quatre dimensions des documents électroniques. *Document numérique, Vol. 1, 1*, 55-60
- Weng, C., Gennari, J., H. (2004) Asynchronous Collaborative Writing through Annotations. *Proceedings of CSCW2004*. Chicago, Illinois, USA, 6-10 November 578-581
- Wolfe, J. (2002) Annotation technologies: A software and research review. *Computers and Composition, 19 (4)*, 471-497

We can work it out: Collaborative Conflict Resolution in Model Versioning[♥]

Petra Brosch*, Martina Seidl, Konrad Wieland, and Manuel Wimmer
Institute of Software Technology, Vienna University of Technology, Austria
lastname@big.tuwien.ac.at

Philip Langer
Department of Telecooperation, Johannes Kepler University, Austria
philip.langer@jku.at

Abstract. For the versioning of code a pantheon of version control system (VCS) solutions has been realized and is successfully applied in practice. Nevertheless, when it comes to merging two different versions of one artifact, the resolution of conflicts poses a major challenge. In standard systems, the developer who performs the later commit is sole in charge of this often time-consuming, error-prone task. This commit carries the inherent danger of losing the modifications of the other developer. Recently, collaborative merge approaches for code versioning systems have been proposed to minimize this risk. In this paper we propose to apply similar techniques in the context of model versioning where the challenge of merging two versions is even more formidable due to their graph-structure and their rich semantics. In particular, modeling is used in the early phases of the software development, where a collaborative merge is beneficial to elaborate a consolidated understanding of a domain.

Introduction

In this paper we describe an extension of the model versioning system AMOR (cf. Altmanninger et al. (2008)) for collaborative conflict resolution. Following this approach, not only one developer is in charge of merging different versions of

[♥]This work has been partly funded by the Austrian Federal Ministry of Transport, Innovation and Technology (BMVIT) and FFG under grant FIT-IT-819584.

*Funding for this research was provided by the fFORTE WIT - Women in Technology Program of the Vienna University of Technology, and the Austrian Federal Ministry of Science and Research.

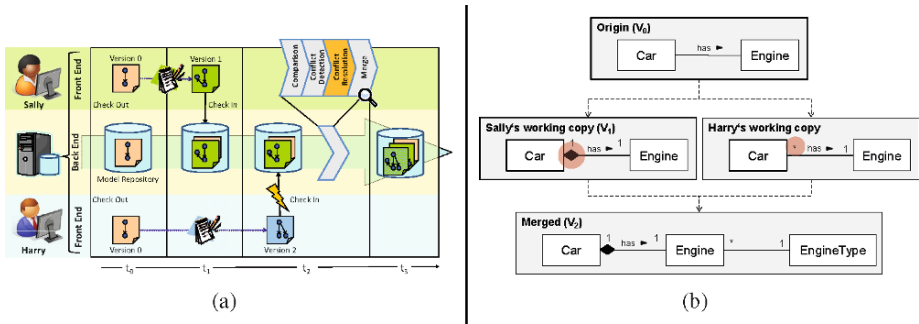


Figure 1. (a) Check-in Process. (b) Motivating Example.

a model but all persons who performed the changes are involved in eliminating the conflicts to obtain one consistent model version.

Collaborative software development without version control systems (VCSs) is nowadays unimaginable. Especially optimistic VCSs are of particular importance because such systems effectively manage concurrent modifications on one artifact performed by multiple developers. In contrast, pessimistic approaches allow only one developer to modify an artifact at the same time. Figure 1(a) illustrates the workflow of applying an optimistic versioning system from the developers' point of view. Two developers (let us call them Harry and Sally) check out the same artifact from a central repository managed by a VCS. When Sally is finished, she loads her new version back to the repository. Later Harry also intends to submit his new version, but the VCS rejects to load his version into the system as his changes are conflicting with the changes of Sally. He is forced to resolve these conflicts before his changes are adopted. The complete responsibility of merging the two versions totally shifts to Harry who may not be aware of Sally's intention and who possibly overwrites or wrongly integrates her modifications. Consequently, the work of Sally could get lost or the merged version could not reflect the intention of Sally. In conventional systems, the consolidation of the other developer is not envisaged. The problems of the asynchronous conflict resolution motivated the augmentation of VCS for code by events notification features as well as with chat possibilities (cf. Fitzpatrick et al. (2006)).

If the artifacts under optimistic version control are text files like source code, merging parallel changes of multiple developers is already a big challenge (cf. Mens (2002)). Merging artifacts as software models is hardly supported by state-of-the-art systems and becomes an even bigger challenge (cf. Barrett et al. (2008); France and Rumpé (2007)). Until now, it was sufficient to version models pessimistically, but as software models nowadays become an indispensable source of information for software engineering—either for documentation purpose or for model driven engineering (MDE) (cf. Bézivin (2005)) where code is automatically generated from models—the need for putting models under optimistic version control arises. Standard VCSs for code usually work on file-level and perform conflict detection by line-oriented text comparison. When applied on the textual serialization of models,

the result is unsatisfactory because the information stemming from the graph-based structure is destroyed and the associated syntactic and semantic information is lost. Consequently, dedicated VCSs for model versioning have been proposed which realize model specific comparison, conflict detection, conflict resolution, and merge components. So far they follow the classical approach where the person performing the later commit of the modified version (i.e., Harry), is left alone with the task of resolving the conflicts. If he has a different understanding of, e.g., the domain, the danger is very high that he destroys the work of the other modeler resulting in unintended models. In this paper we discuss why misunderstandings are even more a potential risk in the context of modeling than in any other area of software development and, inspired by the work of Dewan and Hegde (2007), we propose to integrate a collaborative merging facility into the model versioning system AMOR in order to overcome many problems resulting from the traditional single-person merge process. In particular, modeling is used in the early phases of the software development, where a collaborative merge is beneficial to elaborate a consolidated understanding of a domain.

Motivation

Modeling in parallel and versioning in an optimistic way as already discussed raise several challenges. Concurrent changes of the same model motivated by different but partly overlapping intentions often need manual interventions by the modelers for resolving conflicts. To elaborate on these problems, a concrete example in the context of UML Class Diagram is depicted in Figure 1(b). In this example, after checking out the actual version of the origin model *V0* consisting of the classes *Car* and *Engine* and the association *has*, Sally replaces the association with a composition in her working copy *V1*. Hence, she defines an *Engine* instance as part of one *Car* instance. In parallel, Harry increases the multiplicities in his working copy in a different way (cf. Figure 1(b)) to unbound in order to declare that more than one car may use the same type of an engine (e.g., an engine of the type *Diesel*). Both versions express different understandings of the class *Engine* and therefore an automatic merge is not possible. A naive merge including both variants would lead to a semantically incorrect model as the upper bound for the multiplicity of the composition is restricted to one. A collaborative interaction of both modelers is necessary to find a solution combining both intentions. This exchange of information between the modelers leads to a merged model covering both aspects by introducing a third class named *EngineType* and consequently result in a model of higher semantics and quality. How this collaborative interaction may be integrated in current modeling environments is presented in the next section.

A lot of fertile research effort has been performed to address the challenges of collaborative software development going back to the early nineties (for an overview see Dewan and Riedl (1993)). In a more recent work, Dewan and Hegde (2007) proposed a collaboration model, which enables users to collaboratively resolve conflicts and in further consequence merge their intentions collectively. Still, previous

work only considers the collaborative merge of software code—not software models. An adaption and extension of these concepts in order to match the requirements for collaborative *software modeling* is highly valuable for several reasons.

Software models express aspects of a software system at a much higher level of abstraction and, therefore, reveal a high amount of semantics, domain specific knowledge, and modeling experience. Usually, these skills are distributed variably over all members in a team, which makes the collaboration even more important. The collaboration enables the concentration of these skills—in place and time—where they are needed. Moreover, collaboration is crucial to construct a common vocabulary and understanding. The lack of shared knowledge about goals and concepts of the software often leads to immense non-conformance and low quality in further consequence.

For the sake of minimizing conflicts, development parts are often strictly separated and assigned to team members. Modeling is strongly applied during the requirements engineering phase and analysis phase of a software development project. Since at these early phases the components of the system as well as their borders and interfaces are not clearly defined, the strict separation of tasks is hard to realize. Furthermore, a mistake in early phases leads to disproportionately high costs in later phases. For this reason, it is important to identify design problems as early as possible or even prevent them from the beginning. Design problems in modeling, for example contradicting models, are difficult to detect automatically. The semantic correctness is not expressible in a formal way, so quality assurance mechanisms like unit tests are not applicable. The identification of design problems requires both, modeling as well as domain specific knowledge. Thus, there must be an infrastructure available that supports the joint resolution of conflicts.

Additionally, there exists another requirement for collaborative modeling in contrast to code-centric development: elements as part of a software model may be visualized in more than one diagram. For example, an actor in a UML Use-Case Diagram is often implemented as class within a UML Class Diagram. Both of these two different visual elements (actor and class) represent the same real-world concept. This possibility could recurrently lead to unexpected consequences when concurrent modifications occur and consequently must be considered in conflict detection and resolution.

Finally every member of the team usually feels responsible for his or her work. Using an optimistic versioning approach, coworkers are very often forced to modify a colleague's work in order to resolve a conflict without building a consensus. This frequently leads to semantically inconsistent models, social conflicts and, in the worst case, frustration. Our proposed *collaborative merge* demands to solve such conflicts together, which increases the developers' acceptance as well as the quality of the solution.

We can work it out: Collaborative Conflict Resolution in Model Versioning

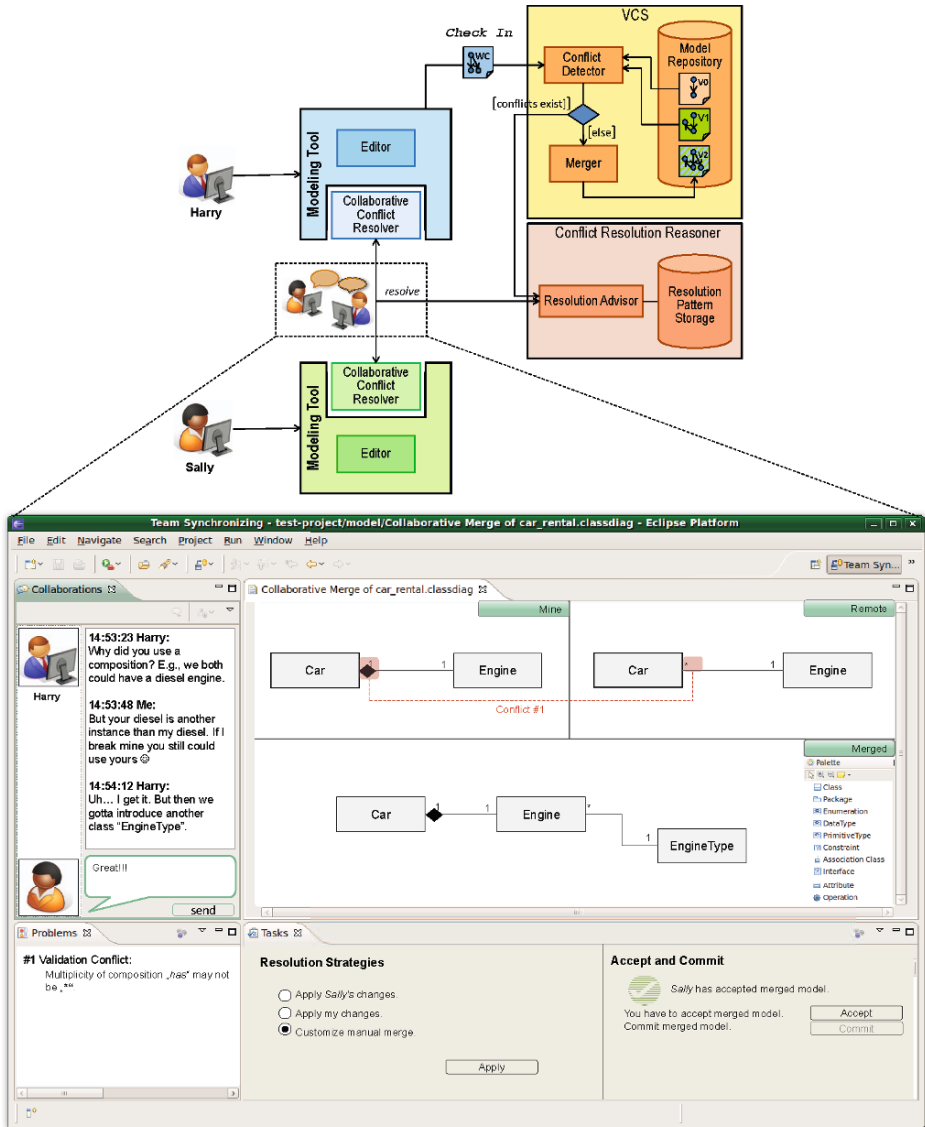


Figure 2. Model Versioning System Architecture and Collaborative Conflict Resolver UI.

Proposed Architecture

In this section, we propose a *semi-collaborative* approach for model versioning. We define semi-collaborative as the parallel editing of artifacts by different modelers is done independently but the task of manual conflict resolution is done in a collaborative manner. In order to overcome the problems mentioned in the previous section, we extend conventional VCSs by a **Collaborative Conflict Resolver** component to support the collaborative resolution of conflicts between model versions. In addition, to provide recommendations in terms of *resolution patterns* for resolving recurring conflicts, we reuse the **Conflict Resolution Reasoner** component of the AMOR VCS (cf. Altmanninger et al. (2008)). The creation of resolution patterns is two-folded. First, a VCS administrator defines a first set of resolution patterns in advance which are then recommended in case such conflicts occur. Second, since it is in general not possible to provide resolution patterns for all kinds of conflicts, the **Conflict Resolution Reasoner** tracks occurred conflicts and their resolution for learning from the user decisions and enhancing the resolution power. For this, the concrete conflict resolutions are broken down to more general facts, formalized in conflict resolution patterns, which are reusable for decision support.

Using the **Conflict Resolution Reasoner** alone, conflicts are semi-automatically resolved when a resolution pattern is applied by one modeler. While this approach already unifies the resolution knowledge of the whole developing team in terms of decision support, it is still not sufficient due to the following reason. The resolution of conflicting changes is completely in the responsibility of the person who checks in later. This procedure neglects information about the intention of the changes of others. Thus, a person should not be left alone with manual conflict resolution, but rather be accompanied by the author of the conflicting version.

For combining the benefits of both approaches, i.e., providing resolution recommendations and a collaborative merge process, we propose an extended VCS architecture as shown in the upper part of Figure 2. The check-in procedure of an extended VCS is explained by means of the previously presented example. As mentioned before, Harry and Sally check out the newest version V0 and start editing. When Sally is finished, she checks in her working copy and luckily she is the first and no conflicts resulting from concurrent changes occurred. Her working copy is persisted in the repository as V1. Later, when Harry has performed his changes, he checks in his working copy. Unfortunately, the **Conflict Detector** reports conflicts and an automatically merge can not be performed. Thus, the conflicts must be resolved manually. The **Resolution Advisor** receives the conflict report and queries the **Resolution Pattern Storage** for recommended resolution patterns. In our scenario, we assume that three recommendations are available, namely, apply only Sally's changes, apply only Harry's changes, or perform a manual merge. No matter which recommendation is followed, a collaborative merge process is necessary, thus, both authors of the two conflicting versions are informed to resolve the conflict together. Both modelers are notified by their **Collaborative Conflict Resolver**, which shows the conflict report and possible resolution strategies.

As illustrated in the lower half of Figure 2, the Collaborative Conflict Resolver shows both versions (cf. Mine window and Remote window) where the conflicting elements are marked, as well as a shared whiteboard for collaborative editing the merged version (cf. Merged window). In this window both modelers are able to change the diagram and elaborate a suitable merged model at the same time. Each modification is immediately visualized on the other modeler's screen. The integrated chat component (cf. Collaboration window), the conflict report (cf. Problems window) and the list of resolution strategies (cf. Task window) help discussing the intention of the changes and finding an appropriate solution, respectively. When both modelers are satisfied with the revised merged version, they each press the *Accept* button and finally the *Commit* button is enabled. Then, by clicking on the *Commit* button the revised merged version V2 is stored in the repository, and the applied conflict resolution is propagated back to the Conflict Resolution Reasoner for learning new resolution patterns. In particular, from this specific conflict and its corresponding resolution, a resolution pattern is mined, namely that in cases there an association is marked as composition and at the same time the multiplicity is set to unbound, an additional class should be introduced. This pattern may be reused for similar examples, such as the following. Consider a model consisting of a class *Library*, a class *Book* and an association between those classes, and concurrently the same modifications as in our running examples occur. One modeler defines that the book is contained in one library, actually meaning with book a concrete book copy, whereas the other defines that a book is offered in several libraries. By applying the previously explored resolution pattern, an additional class *BookCopy*—the name has to be inserted by the modelers—is introduced in order to resolve the contradicting association definition.

Challenges

We integrate the Collaborative Conflict Resolver in the VCS AMOR (cf. Altmaninger et al. (2008)) which provides not only dedicated versioning support for models, but also an advanced conflict resolution component. With the power of these elements, we expect to improve the check-in process in two ways. First, the quality of the merged version increases through the consolidation of all parties involved in modifications. Second, the double manpower for the conflict resolution combined with the solution strategies offered by the Conflict Resolution Reasoner of AMOR may decrease the time spent on model merging. We are aware of the multiple challenges on the way to the realization and practical application. In the following, we shortly discuss some of our concerns and outline solution statements.

User Acceptance. The success of our component highly depends on the willingness of modelers to communicate with each other and to engage with unfamiliar elements newly introduced in the versioning process. We expect a high acceptance by the users as soon as the advantages become evident, i.e., they save time and can focus on their own work and not on the work of others.

Visualization. Strongly connected to the user acceptance is the provision of an attractive user interface. Of paramount importance is the well-arranged presentation of the very diverging information. In the previous section we have already presented a prototypical depiction of the **Collaborative Conflict Resolver**, but we aware that if our system should be used in the large scale, much effort will have to be put in the evaluation and the resulting suggestions.

Scalability. In the previous scenario only two modelers were involved in the conflict situation. In practice, even more persons will work on one model and quite naturally, the following question arises. How does the system behave if the modification of more than two developers are conflicting. In practice, this does not yield the problem, as it is possible to trace back the situation to the base case of only two conflicting versions, because a working copy only has to be merged with the actual version in the repository.

Evaluation. We plan to conduct case studies where we evaluate the fertility of our approach. As our system highly depends on user interaction, tests which are not performed in the context of a real world project of reasonable size (more than one modeler) are not expressive. We have the opportunity to assess the viability of our new component with students of a university course as well as in the context of a real world test bed provided by our industrial partner Sparx Systems, the vendor of Enterprise Architect™.

References

- Altmanninger, K., G. Kappel, A. Kusel, W. Retschitzegger, M. Seidl, W. Schwinger, and M. Wimmer (2008): ‘AMOR—Towards Adaptable Model Versioning’. In: *Proc. of the 1st International Workshop on Model Co-Evolution and Consistency Management*.
- Barrett, S., P. Chalin, and G. Butler (2008): ‘Model Merging Falls Short of Software Engineering Needs’. In: *Proc. of the 2nd Workshop on Model-Driven Software Evolution*.
- Bézivin, J. (2005): ‘On the Unification Power of Models’. *Journal on Software and Systems Modeling*, vol. 4, no. 2, pp. 171–188.
- Dewan, P. and R. Hegde (2007): ‘Semi-synchronous conflict detection and resolution in asynchronous software development’. In: *Proc. of the 10th European Conference on Computer-Supported Cooperative Work*. pp. 159–178, Springer.
- Dewan, P. and J. Riedl (1993): ‘Toward Computer-Supported Concurrent Software Engineering’. *IEEE Computer*, vol. 26, no. 1, pp. 17–27.
- Fitzpatrick, G., P. Marshall, and A. Phillips (2006): ‘CVS integration with notification and chat: lightweight software team collaboration’. In: *Proc. of the 2006 ACM Conference on Computer Supported Cooperative Work*. pp. 49–58, ACM.
- France, R. B. and B. Rumpe (2007): ‘Model-driven Development of Complex Software: A Research Roadmap’. In: *Proc. of the 29th International Conference on Software Engineering*. pp. 37–54, IEEE Computer Society.
- Mens, T. (2002): ‘A State-of-the-Art Survey on Software Merging’. *IEEE Transactions on Software Engineering*, vol. 28, no. 5, pp. 449–462.

On the effects of Refactoring in the Coordination of Software Development Activities

Cleidson R. B. de Souza, Maryanne P. Rosa, Crys S. Goto, Jean M. R.
Costa, Pedro J. F. Treccani

Faculdade de Computação, Universidade Federal do Pará,

Belém, PA, Brasil

cdesouza@ufpa.br

Abstract. *Several empirical studies suggest that an alignment between the architecture of a software system and the coordination of development activities lead to better quality and improved performance. In this paper we investigate the possible effects of misalignments due to changes in the software architecture by describing the results of an exploratory study about the effects of refactoring in the coordination of software development activities in an open source project. We studied refactorings because they are perfect examples of changes in the software architecture. The project evaluated is the Jackrabbit, an Apache Software Foundation project. This project was analyzed using statistical tests and social networks analysis metrics. We evaluate different hypothesis regarding the impact of the refactoring process on project coordination. Initial results suggest that core software developers are especially affected by refactoring activities.*

Introduction

More than forty years ago, Conway (1968) suggested that the relationship between the architecture of a software system and the structure of the organizational developing this software is homomorphic, the Conway's Law. Parnas (1972) similarly suggested a mechanism, the information hiding principle, to structure the software architecture in order to reduce software developers' coordination needs. Recently, these theoretical proposals have been corroborated by several qualitative (Staudenmayer 1997; Grinter 1999; de Souza et al. 2004; de Souza and Redmiles 2008) and quantitative (Sosa et al.

2004; Cataldo et al. 2006; Cataldo 2007) empirical studies. In sum, both researchers and practitioners recognize that the communication and coordination effort necessary to develop a software system is closely linked to its architecture. In fact, researchers have studied the alignment between the architecture and the coordination. For instance, software developers' performance is related to how well software developers align their coordination efforts with the dependencies in the software architecture, both at the team level (Staudenmayer 1997), and at the individual level (Cataldo 2007). On the other hand, misalignment between these aspects is seen as a possible explanation for breakdowns in software development projects (Brooks 1974; Bass et al. 2007). However, these studies are limited because they do not isolate the consequences of this misalignment: there are confounding factors.

In this paper we aim to investigate, in isolation, the possible effects of the misalignment between the software architecture and the coordination. Our approach in this paper is to study the effects of changes in the software architecture into software developers' coordination, since, in order to be successful when one aspect changes, the other must change accordingly (Staudenmayer 1997; Cataldo et al. 2006). We adopt a simple and informal definition of coordination based on the analysis of the work performed by software developers in their mailing lists. We report our results based on an analysis of a case study of software refactoring in an open source project. Software refactoring is the process of changing a software system in such a way that does not alter the external behavior, yet improves its internal structure (Fowler et al. 1999), therefore, it adequately represents changes in the software architecture. The coordination of the software development work is analyzed using statistical tests and social network analysis methods (Wasserman and Faust 1994). Our initial results suggest that key project members are affected by the coordination, since they perform additional work during the refactorings.

The rest of this paper is organized as follows. Our methods are described in the following section. Then, our results and discussion about them are presented. After that, we present our conclusions and future work.

Methodology

The research method used in this paper is an exploratory, single-case study (Yin, 2004). In any given case-study, it is necessary to detail how the case was selected. This is detailed in the following section.

Case Selection

In order to understand the effects of refactoring, we decided to analyze an open source project, i.e., a case is a software project. The project used is the Jackrabbit¹, an Apache

¹ <http://jackrabbit.apache.org/>

Software Foundation project. This project was selected because its information is publicly available in the internet, and, more importantly, it has had several refactorings since its inception.

This project has four different mailing lists: announcements (Jackrabbit Announce List), users (Jackrabbit Users List), development (Jackrabbit Development List), and source control (Jackrabbit Source Control List). In this work, we focused solely on the development list since we are interested in the coordination of *software developers*, the announcement list is used to inform users about new releases, and the source control list contains messages automatically generated by the configuration management tool.

Once our project was selected, we proceeded to collect and analyze its data as described in the following section.

Data Collection and Analysis

In order to collect data about the jackrabbit project we used a web-crawler named OSSNetwork (Balieiro et al. 2008) to extract information from the jackrabbit web-site. This tool allows one to: retrieve information from FLOSS repositories, store this information in a database, generate different types of social networks from this information, and, finally, analyze these networks.

Using OSSNetwork, we collected data about the development mailing list and stored it on a relational database. Queries to the database were performed to extract information like: thread duration (in days), total number of messages per thread, and number of different developers involved in each thread.

Data was analyzed in two ways: using statistical tests (Wild and Seber 1999) and using social network analysis methods (Wasserman and Faust 1994). Social networks were created based on the information extracted from the mailing list: software developers were the nodes of the network while edges were created when a developer, let's say B, replied to another developer, e.g. A. In this case, an edge from B to A is created (Wasserman and Faust 1994).

Identification of the Refactorings

We identified the refactorings in the project by reading messages in the mailing list and by reading the description within the bug-tracking system used, JIRA. We basically searched for the word "refactoring" in the mailing list and bug-tracking system. In the jackrabbit project, each refactoring is mapped to a JIRA issue so that there is a HTML page for each issue / refactoring containing information about the beginning and end date of the refactoring, the files modified, its author, among other types of information.

We identified a total of eight (8) refactorings in the project. Some of them happened in a short period of time (e.g., one day). We decided not to analyze these refactorings because we wanted to analyze the effects of refactoring while it was taking place (more details in the next section). In this paper we report the preliminary results from our analysis of two refactorings:

- Refactoring 1 (R1) – it took place between March 03, 2005 and March 16, 2005. This refactoring can be mapped to JCR-53 and JCR-66 in the JIRA tool; and
- Refactoring 2 (R2) – it took place between Jan 256, 2007 and March 07, 2006 and is mapped to issue JCR-309.

Once we defined the refactorings to be analyzed, it was necessary to define the periods of analysis for each refactoring.

Periods of Analysis

Refactorings R1 and R2 were analyzed according to two periods: *before* and *during* the refactorings. *Before* a refactoring means the period of one month before the refactoring started, while *during* a refactoring means the period in which the refactoring took place (start and end date of the refactoring according to the JIRA system). The period before refactoring R1 (R2) is labeled B1 (B2), while the period during refactoring R1 (R2) is called D1 (D2).

For each period of each refactoring, we calculated the mean and the standard deviation for each thread in the mailing list, the number of messages per thread and the total number of different developers who posted messages in the thread. If a developer posted more than one message in the same thread, he is counted as only one developer, but all his messages are taken into account.

Tests performed

We initially performed statistical tests to verify whether the data collected – thread duration, number of messages per thread and number of collaborators per thread – was normally distributed. Since this was not the case, we used non-parametrical tests (Wild and Seber 1999).

In our first test, we looked at thread duration, number of messages and number of collaborators per thread using the Mann-Whitney U test (Wild and Seber 1999). This test aims to compare independent samples for assessing whether two samples come from the same distribution. We compared periods B1 and D1, and B2 and D2 to find out information about these periods. Therefore, the Mann-Whitney U test will compare whether the duration of the threads before (B1) a refactoring is different from the duration of the threads during (D1) the same refactorings.

The second test aimed to compare the number of messages sent by each developer before and during each refactoring. In this case, we used the Wilcoxon signed-rank test, a non-parametric statistical test for two related samples or repeated measurements on a single sample (Wild and Seber 1999). Samples are related in this case because they describe the number of messages sent by the same developer before and during a refactoring. In this case, this test will evaluate whether the same developer sends a similar number of messages before and during a refactoring.

Finally, we tested the number of messages sent by a specific group of developers: the developers who were in the core (Borgatti 2000) of (i) the social network of the period before the refactoring, *and* of (ii) the social network created for the period during the refactoring. This test was performed only for those developers who were present in both social networks and who were in the core of each network to identify the impact caused by the refactoring in the most important developers of the project. Network core was identified using UCINET², a traditional social network analysis tool.

Results

As mentioned in the previous section, in our first test we looked at thread duration, number of messages and number of collaborators per thread. In this case, we found no statistical difference between periods before and during refactoring for refactoring R1 ($p=0,647$ for duration, $p=0,786$ for the number of messages and $p=0,636$ for the number of developers per thread), but a significant value for refactoring R2 ($p=0,004$ for duration, $p=0,016$ for the number of messages and $p=0,006$ for the number of developers per thread).

We tried to find a possible explanation for this result – a significant value in R2, but no significance in R1 – based on a qualitative analysis of these refactorings. We observed that refactoring R2 caused API changes across several modules, i.e., changes with broader impact than the other refactoring (R1), which changes a single module. To be more precise, the exact description of refactoring R2 extracted from the JIRA bug-system is:

“To better document and track the public JCR extensions and component API provided by Jackrabbit and to allow more room for refactoring within the Jackrabbit core, we should move (or create) the supported API interfaces to a new org.apache.jackrabbit.api package.”

In fact, by carefully examining the JCRs, we identified that R1 caused 110 changes in the source code, while refactoring R2 lead to 420 changes.

We also tested the average number of messages sent by each developer who posted messages in the mailing list during the analyzed periods. We found out that developers, during the periods of refactoring, send more messages than in the periods that precede a refactoring. This is true for both refactorings R1 and R2. To be more precise, we found some evidence of this result in R1 ($p=0,055$) and strong evidence in R2 ($p=0,0001$). This is an interesting result because it suggests that software developers are performing additional work during the refactorings.

However, this might not be a problematic situation per se because those performing additional work might be less important developers doing work in the periphery of the project. Therefore, we decided to perform the same test, but now considering only those

² <http://www.analytictech.com/ucinet/ucinet.htm>

developers who were in the core of the project, the key project members. In order to do that, we used social network analysis to identify the developers in the core and in the periphery of the project in both periods (before and during a refactoring) and for both refactorings (R1 and R2). We then used the same statistical test as before (Wilcoxon signed-rank) to evaluate the number of messages sent per developer, but now considering only those software developers who were in the core of the social network in both periods. We found out that developers in the core of the projects also send more messages during a refactoring. However, now we have very strong evidence supporting this result in both refactorings (R1 with $p=0,002$ and R2 with $p=0,004$).

Discussion

Our results suggest that the coordination of software development projects is in fact affected by refactorings of the source code. We initially evaluated discussion threads in the jackrabbit development mailing list and observed they were not affected, regarding its number of messages, duration or number of software developers involved for the refactorings analyzed, by the refactorings. This seems to suggest that the project community as a whole does not performed differently during the refactoring.

However, the results of our preliminary analysis indicate that some software developers send much more messages during the refactoring than in the period preceding it. More importantly, those software developers are the key members of the project, those who belong to the core of the social network created from the mailing list: they are the ones who are more active in the list. In general, we argue that we were able to find out novel and important effects of the misalignment of the software architecture and the coordination.

It is interesting to note two aspects: the overall number of messages per thread does not increase during refactoring, although the number of messages per contributor does change. We investigated this counter-intuitive result and observed that some contributors who *did not* participate in the discussions before the refactoring, *did* participate in the discussion during the refactoring. A possible explanation is that those developers were aware of the refactoring and decided to “jump in” the discussions to contribute. Additional research is still necessary to investigate alternative explanations.

In general, a consequence of our results is the recognition that changes in the software architecture need to carefully planned to not disturb even further core developers of the project. For instance, before a refactoring takes place, one needs to take into account its possible implications: core developers will need to prepare themselves to potentially guide other developers during the process. In addition, changes in the code need to be designed in order to avoid additional coordination problems. In fact, after performing the statistical tests, we analyzed the messages being exchanged during the refactorings. We were able to identify, at least, three situations in which software developers performed conflicting changes in the code, which lead to additional communication and coordination efforts.

Conclusions and Future Work

In this paper we describe our research on the effects of refactoring in the coordination of software development activities. This research is motivated by theoretical and empirical studies (Conway 1968; Parnas 1972; de Souza et al. 2004; Cataldo et al. 2006) regarding the socio-technical relationship between software architecture and coordination of development activities. Based on this socio-technical relationship, the hypothesis of the presented work is that refactoring, as it modifies the structure of the architecture, has effects in the coordination of software development activities. In this paper, we test this hypothesis with an exploratory single-case study, whose goal is to examine these effects in an open source project.

We analyzed the development mailing list of an open source project called jackrabbit using non-parametrical statistical tests and social network analysis and found out that while a refactoring is being performed, software developers who are central to the project engage in *additional* communication with their colleagues. We are not able to classify this effect as either negative or positive, but these results suggest the need to carefully plan refactorings in order to avoid an overload to core software developers in the project. It is important to emphasize that this paper reports on a single-case study in an open source project, therefore further research is necessary to generalize the results.

As for future work, we plan to analyze additional refactorings and projects in order to better understand the socio-technical relationship between software architecture and coordination, and its effects. We also plan to use qualitative methods to analyze the content of the messages exchanged during the refactorings.

Acknowledgments

This research was supported by the Brazilian Government under grants CNPq 479206/2006-6, CNPq 473220/2008-3 and by the Fundação de Amparo à Pesquisa do Estado do Pará (FAPESPA) through "Edital Universal N.º 003/2008".

References

- Balieiro, M. A., et al. (2008). Facilitating Social Network Studies of FLOSS using the OSSNetwork. *International Conference on Open Source Systems*. Milan, Italy, Springer Series in Computer Science: 343-350.
- Bass, M., et al. (2007). Architectural Misalignment: An Experience Report. *IEEE/IFIP Working Conference on Software Architecture*.
- Borgatti, S. (2000). "Models of Core-Periphery Structures." *Journal of Social Network Analysis*. **21**: 375-395.
- Brooks, F. P. (1974). *The Mythical Man-Month: Essays on Software Engineering*, Addison-Wesley.
- Cataldo, M. (2007). *Dependencies in Geographically Distributed Software Development: Overcoming the Limits of Modularity*. School of Computer Science. Pittsburgh, PA, Carnegie Mellon University. Ph.D.: 188.

- Cataldo, M., et al. (2006). Identification of Coordination Requirements: implications for the Design of Collaboration and Awareness Tools. *20th Conference on Computer Supported Cooperative Work*. Banff, Alberta, Canada, ACM Press.
- Conway, M. E. (1968). "How Do Committees invent?" *Datamation* **14**(4): 28-31.
- de Souza, C. R. B. and D. Redmiles (2008). An Empirical Study of Software Developers' Management of Dependencies and Changes. *International Conference on Software Engineering*. Leipzig, Germany, IEEE Press.
- de Souza, C. R. B., et al. (2004). How a Good Software Practice thwarts Collaboration - The Multiple roles of APIs in Software Development. *Foundations of Software Engineering*, Newport Beach, CA, USA, ACM Press.
- Fowler, M., et al. (1999). *Refactoring: Improving the Design of Existing Code*, Addison-Wesley Professional.
- Grinter, R. E. (1999). System Architecture: Product Designing and Social Engineering. *Work Activities Coordination and Collaboration*, San Francisco, CA, USA, ACM Press.
- Parnas, D. L. (1972). "On the Criteria to be Used in Decomposing Systems into Modules." *Communications of the ACM*. **15**(12): 1053-1058.
- Sosa, M. E., et al. (2004). The Misalignment of Product Architecture and Organizational Structure in Complex Product Development. *Management Science*. **50**(12): 1674-1689.
- Staudenmayer, N. A. (1997). Managing Multiple Interdependencies in Large Scale Software Development Projects. *Sloan School of Management*. Cambridge, MA, USA, Massachusetts Institute of Technology. Ph. D.
- Yin, R. (2004). *Case study research: Design and methods*. Beverly Hills, CA: Sage Publishing.
- Wasserman, S. and K. Faust (1994). *Social Network Analysis: Methods and Applications*. Cambridge, UK, Cambridge University Press.
- Wild, C. J. and G. A. F. Seber (1999). *Chance Encounters: A First Course in Data Analysis and Inference*, John Wiley & Sons.

Divided by a common acronym: *On the fragmentation of CSCW*

Kjeld Schmidt

Copenhagen Business School

schmidt@cscw.dk

Abstract. CSCW is in an advanced state of fragmentation. The acronym now, by and large, denotes widely diverging research programs that, apart from a shared name, have little or nothing in common. This situation obviously calls for clarification. Recounting the prehistory and formation of CSCW, the paper shows that CSCW, as a distinct research program devoted to the development of new technologies on the basis of understanding actual cooperative work practices, arose in response to the crises in which ‘Computer Mediated Communication’ (CMC) and ‘Office Automation’ (OA) had landed by the late 1980s. The paper finally discusses the reasons why CMC, although superseded as a research paradigm by the practice-oriented program of CSCW, has gained a new lease on life in CSCW and thus why CSCW has become fragmented.

Development of technology results in technical knowledge, methods, principles, etc. That is, it is essentially a conceptual effort. CSCW research therefore has to be cumulative. For CSCW research to be cumulative does not mean a linear process, of course, but a process in which the different contributions — empirical studies, conceptualizations, experimental designs, architectures — build upon, corroborate, exemplify, complement, generalize, question, discuss, subvert, or overthrow other contributions. However, in a ‘degenerating’ research program, to use the term coined by the philosopher of science Imre Lakatos (1970), this process is effectively blocked. The required continuity of the program, the ongoing development of concepts and frameworks, is replaced by restless reformulation of the research problem, slapdash changes of scope, unaccountable redefinitions of key concepts, etc. Under such circumstances the progressive development of the conceptual foundation of technology is not possible.

Now, CSCW was never a well-defined research area. In 1988 Liam Bannon *et al.* observed that ‘for the moment the name CSCW simply serves as a useful forum for a variety of researchers with different backgrounds and techniques to discuss their work’ (Bannon, *et al.*, 1988). But as indicated by the temporal modifier ‘for the moment’, Bannon *et al.* obviously expected this state of affairs to pass. Others, however, saw in this condition a virtue of the field: ‘Perhaps paradoxically, one of the most refreshing things about CSCW may be the fact that its meanings [are] debated. For as long as this is the case, researchers will reflect upon the nature of their work, what its aims and outcomes are or might be’ (Bowers and Benford, 1991, p. 1). Neither of these expectations have been fulfilled. In fact, CSCW has become fragmented. The upshot is that CSCW as a research field is unable to proceed in anything like a cumulative fashion and thus, generally, unable to contribute systematically and constructively to the development of new technologies.

One example will suffice to indicate the level of fragmentation. Take the review article in the *HCI Handbook*, entitled ‘Groupware and computer-supported cooperative work’ (Olson and Olson, 2003). Written by two eminent CSCW veterans, Gary Olson and Judith Olson, the article gives an overview of a range of types of application such as email, conferencing, instant messaging, group calendars, shared repositories and work spaces, media spaces, and collaborative virtual environments. Although the authors assert to be using the terms groupware and CSCW ‘quite broadly’ (p. 584), they nonetheless conceive of the field quite narrowly as a field focusing on technologies that, in different ways, support ‘collaboration’ or ‘group work’ ‘at a distance’ (*ibid.*).

This account is puzzling. First of all, it assimilates CSCW with the research area of Computer-Mediated Communication (CMC) that preceded it by many years. But what makes it quite remarkable is that it leaves out the significant body of CSCW research that involves investigations of cooperative work practices in professional settings (such as air traffic control, maintenance work, urban rapid transit control, software engineering, manufacturing, health care). This entire body of research is not reflected in the review at all. In fact, this review of CSCW completely ignores the substantial contribution of ethnographic or workplace studies to CSCW. Terms such as ‘ethnography’, ‘workplace study’, or ‘field work’ do not even feature in the article, nor are there any references to particular ethnographic studies or to the large CSCW literature about the role of workplace studies.

This is remarkable because ethnographically informed CSCW research has had deep and lasting impact on the scope and direction of major sectors of CSCW due to the way in which in-depth studies and sociological analysis of actual cooperative work practices have been made to bear on conceptual issues of technological research and development. Furthermore, exactly by virtue of this paradigmatic kind of socio-technical inquiry, CSCW has exerted significant influence on re-

lated scholarly fields such as human-computer interaction (HCI), participatory design (PD), and science and technology studies (STS).

As a researcher who, for decades, has been involved in ethnographically informed CSCW research, I was, at first, appalled at a review of CSCW so conspicuously partial. But then again, the picture painted by Olson and Olson is also, in a sense, a genuine reflection of the actual state of affairs in CSCW, and I could easily imagine another review article, from another quarter, that was equally partial. In that light, moral outrage is unwarranted. What we have, it seems, is rather a situation characterized by ‘incommensurate paradigms’ (Kuhn, 1962): different research programs that have little or nothing in common, apart from a shared acronym, addressing different problems, conceived of in different conceptual frameworks, employing largely disparate methods and techniques of research. But what we have is not a classical paradigm crisis. The mutual indifference of the different programs in CSCW is rather an indication of fragmentation.

This situation calls for clarification of the distinct features of the different research programs. The present paper is meant as a contribution to this process. There are other fault lines and other sources of fragmentation in CSCW, but for obvious reasons the paper will focus on showing the fundamental differences between the research program represented in the review article by Olson and Olson and the practice-oriented program. An initial look at the prehistory and formation of CSCW will show that the practice-oriented program of CSCW emerged in response to internal problems in CMC as it was then conceived and in other contemporary research areas and will thereby help to clarify what was and is specific in CSCW. The purpose of this is not to write a history of CMC and CSCW but to demarcate the intellectual fault lines.

1. The prehistory of CSCW

The beginning of CSCW is so humble that there hardly is any record of it: small practical steps to deal with mundane practical problems.

We are back in the prehistory of CSCW, in the early days of electronic computers, when the notion of computer-mediated communication was gradually gestated. More precisely, the notion of CMC begins with the notion of ‘time-sharing’ operating systems that matured around 1960. Computer systems were at that time excessively expensive and it was mandatory that systems were operating close to full capacity. Consequently, the few computers that were around were running in a batch-processing mode, one job after another on a ‘first-in, first-served’ basis, or, as it was aptly expressed by J. C. R. Licklider, who played a leading role in the early development of interactive computing, the ‘conventional computer-center mode of operation’ was ‘patterned after the neighborhood dry cleaner (“in by ten, out by five”)’ (Licklider and Clark, 1962, p. 114). This economic regime effectively precluded computer applications such as military command and control,

war gaming, air traffic control, computer-aided design, etc. that were of obvious and critical interest to important stakeholders such as the US military and other branches of government. The same ‘in by ten, out by five’ regime also made programming, especially debugging, a deadening affair. This gave ordinary computer technicians a strong motive for devising alternative modes of operation. So, around 1960 the idea of letting a central computer system service several users ‘simultaneously’ was hatched. In the words of John McCarthy, one of the fathers of the idea, the solution was an operating system that would give ‘each user continuous access to the machine’ and permit each user ‘to behave as though he were in sole control of a computer’ (McCarthy, 1983). The first running operating system of this kind seems to have been the Compatible Time-Sharing System or CTSS which was launched in 1961. The various users were connected to the ‘host’ computer via terminals and each would have access to the computing power of the ‘host’ as if he or she was the only user.

Now, the users of the first of these systems were typically engaged in cooperative work. Some were engaged in developing operating systems or other large-scale software projects and were, as a vital aspect of this, engaged in various forms of discourse with colleagues within the same project teams and research institutions, that is, with colleagues already connected to the central computer system. Likewise, software technicians would need to coordinate with system operators about possibly lost files to be retrieved, about eagerly-awaited print jobs in the queue, etc. The time-sharing operating system they were building or using provided a potential solution to this need, and the idea of using the system to transfer text messages from one worker to another did not require excessive technical imagination. As one of the designers of one of the first email systems recalls:

‘[CTSS] allowed multiple users to log into the [IBM] 7094 from remote dial-in terminals[] and to store files online on disk. This new ability encouraged users to share information in new ways. When CTSS users wanted to pass messages to each other, they sometimes created files with names like TO TOM and put them in "common file" directories, e.g. M1416 CMFL03. The recipient could log into CTSS later, from any terminal, and look for the file, and print it out if it was there.’ (Van Vleck, 2001)

A proper mail program, ‘a general facility that let any user send text messages to any other, with any content’ was written for CTSS by Tom Van Vleck and Noel Morris in the summer of 1965 (*ibid.*). It allowed one programmer to send a message to individual programmers, provided one knew the project they worked on, or to everybody on the same project. The message was not strictly speaking ‘sent’; it was appended to a file called MAIL BOX in the recipient’s home directory. The same year Van Vleck and Morris also devised a program (*.SAVED*) ‘that allowed users to send lines of text to other logged-in users’, that is, a primitive form of ‘instant messaging’ (*ibid.*).

The scope of the exchange of messages with these and similar programs was limited by the boundary of the hierarchy comprising the local central computer

system and the terminals connected to it. Messages could not travel beyond the event horizon of this black hole.

This world of isolated systems dissolved with the development of network computing. Again the motivation driving the development was not to develop facilities for human interaction, not to mention cooperative work, but to utilize scarce resources in a more economical way. As pointed out by Ian Hardy, in his very informative history of the origins of network email,

‘ARPANET planners never considered email a viable network application. [They] focused on building a network for sharing the kinds of technical resources they believed computer researchers on interactive systems would find most useful for their work: programming libraries, research data, remote procedure calls, and unique software packages available only on specific systems.’ (Hardy, 1996, p. 6).

For Licklider, who also initially headed the development of ARPANET, the motivation for the network was to reduce ‘the cost of the gigantic memories and the sophisticated programs’. When connected to a network, the cost of such shared resources could be ‘divided by the number of users’ (Licklider, 1960). That is, the primary motive was again economic.

Anyway, after pioneering work on the underlying packet-switching architecture and protocols, the experimental ARPANET was launched in 1969, connecting measly four nodes. In the summer of 1971, when the network had expanded to fifteen nodes, a programmer named Ray Tomlinson devised a program for sending email over the network. He recalls that, while he was making improvements to a single-host email program (`SNDMSG`) for a new time-sharing operating system, ‘the idea occurred to [him]’ to combine `SNDMSG` with an experimental file-transfer protocol (`CPYNET`) to enable it to send a message across the network, from one host to another, and append it to the recipient’s `MAILBOX` file. An instant success within the tiny world of ARPANET programmers, this very first network email program triggered a chain reaction of innovation that within less than a couple of years resulted in the email designs we use today: a list of available messages indexed by subject and date, a uniform interface to the handling of sent and received mail, forwarding, reply, etc. — all as a result of programmers’ improving on a tool they used themselves. Within five years or so, the volume of email messages had become one of the heaviest traffic component on the growing network (Hardy, 1996, p. 21), and in 1977 an official ARPANET standard for electronic mail was adopted (Crocker, *et al.*, 1977).

What is remarkable in this story, and what also surprised those involved when they began to reflect on the experience, was ‘the unplanned, unanticipated and unsupported nature of its birth and early growth. It just happened. and its early history has seemed more like the discovery of a natural phenomenon than the deliberate development of new technology’ (Myer and Dodds, 1976, p. 145). And at a meeting in January 1979, convened to discuss the ‘the state of computer mail in the ARPA community and to reach some conclusions to guide the further development of computer mail systems’, it was ‘noted’ as a fact ‘that most of the mail

systems were not formal projects (in the sense of explicitly sponsored research), but things that “just happened” (Postel, 1982, p. 2). The history of network email after that is well known. The technology migrated beyond the small community of technicians engaged in building computer networks to computer research in general and from there to the world of science and eventually to the world at large.

That is, as in the case of local email on time-sharing operating systems, network email came as an afterthought, devised by computer technicians for their own use, as a means for coordinating their cooperative effort of building, operating, and maintaining a large-scale construction, in this case the incipient Internet. This pattern would repeat itself, again and again. Email and many other CMC technologies that came later were typically thrown together like the scaffolding at a construction site only to become a main feature, relegating the resulting building itself, which had been the original and official objective, to something close to a support structure (cf. also Gillies and Cailliau, 2000).

1.1. The rise and fall of CMC

The experience that human interaction could be facilitated by computers, as demonstrated by email and other protocols, immediately caught the attention and imagination of technologists, who then enthusiastically began developing a generalization of the message exchange idea underlying email, which was soon dubbed ‘computer conferencing’ (for an overview of this work, cf. Kerr and Hiltz, 1982). In its simplest realizations ‘computer conferencing’, in contrast to email, was not restricted to point-to-point message exchanges but supported public exchanges within the forum of the online ‘conference’, regulated in accordance with some established structure. ‘Conferencing’ was in fact often advocated as a remedy for the ‘information overload’ which was seen as an inexorable consequence of point-to-point message exchange (Palme, 1984; Hiltz and Turoff, 1985). The more ambitious experiments, such as EMISARI and EIES (e.g., Turoff, 1972, 1973) and FORUM (Vallee, 1976), explored the rather grand design vision of group communication structured according to some presumptively rational model. Sometimes the experiments allowed for long-term use and thus evolution of ‘user behavior’ (e.g., Hiltz and Turoff, 1981).

While not a development activity undertaken by technicians for their own benefit, this line of research was still characterized by relatively close coupling of experimental design and evaluation work. For instance, between 1973 and 1975, FORUM was tested in 28 conferences and improvements ‘were rapidly incorporated’ (Vallee, 1976; Panko, 1977).

Although the experiments with ‘computer conferencing’ at the time were seen as very promising and reported as very successful, this particular research program ran out of steam. This has to do with underlying conceptual limitations. ‘Computer conferencing’ research shared with the standard message exchange paradigm the presumption that human communication generally is or can be

treated as a distinct activity. True, workers do interrupt their primary work to have conversations and exchange notes, letters, memos about their work (and about other matters). They also, occasionally, put their work aside to go to meetings. For some workers, e.g., managers, the major part of their work day may be spent in conversations and meetings. But apart from managerial work and in the greater scheme of things, conversations and meetings are exceptions, interruptions, ‘a necessary evil’ perhaps, or simply considered ‘a waste of time’. And even when workers engage in conversations and meetings, such discourses are generally related to the state of affairs in their work, to the flow of work, the schedule, the production facilities, the archives, and in their deliberations workers will discuss schedules, plans, schemes, and so on; they will collate, arrange, distribute, present, hand out, walk up to, gather around, point to, gesture at, inspect, amend, etc. all sorts of artifacts.

By the mid-1980s this insight began to mature and be voiced (cf., e.g., Bannon, 1986, p. 443). The CMC research program had landed in a crisis.

The critique of the underpinnings of CMC was expressed clearly and succinctly by Irene Greif in her ‘Overview’ of CSCW in her influential *CSCW: A Book of Readings* (1988). Having noted the rapid development of CMC from electronic mail to computer conferencing she then observes:

‘Computer conferencing has since been expanded to support a wide range of “many-to-many communication” patterns. However, when computer conferencing is applied to some task, the model breaks down. The unstructured body of messages is suitable for the free-flowing text of natural language, but does not let us set the computer to work on our problems. Designers who draw pictures, software developers who jointly write code, financial analysts who collaborate on a budget — they all need coordination capabilities as an integral part of their work tools. That means coordination support within the CAD engineer’s graphics package, within the programmer’s source-code editor, within the budget writer’s spreadsheet program. It means support for managing versions of objects, be they pictures, programs, or spreadsheets. It means ways to distribute parts of the object for work by contributing group members, ways to track the status of those distributed parts, ways to pull completed objects back together again. The limit of electronic mail and computer conferencing is that they have such features for managing messages only. CSCW widens the technology’s scope of application to all the objects we deal with.’ (Greif, 1988, pp. 7 f.)

Greif’s judgment that ‘the model breaks down’ was mirrored in the European CMC research community. This community had emerged in the wake of the European efforts to develop computer networking (cf. Gillies and Cailliau, 2000). As TCP/IP slowly became available in operating systems and developers began to be able to take it for granted, and as the ‘message handling’ standards stabilized in the first half of the 1980s (X.25, X.400, STML), European CMC researchers such as Rolf Speth, Uta Pankoke-Babatz, Wolfgang Prinz, Steve Benford, and others, organized in the AMIGO project, embarked on what was seen as the logical next step, namely, developing the standards required for putting it all together: email as well as directories, calendars, schedules, and so on.

However, the European CMC researchers soon realized that the ‘message-handling’ model underlying CMC was too limited (Pankoke-Babatz, 1989). In work practices, communication is normally not a separate activity; it is typically an integrated aspect of doing the work. In fact, exchanging messages usually presumes that work is interrupted. It was therefore considered necessary to be able to incorporate communication functionality in the various domain-specific applications.

On the other hand, the European CMC researchers rejected the ‘computer conferencing’ paradigm as a way to provide structure to the exchange of messaging. Guided by ‘a strong commitment to the actual situation in working life’ (Pankoke-Babatz, 1989, p. 20), they repudiated the idea underlying the ‘computer conferencing’ paradigm of providing ‘a new model’ of communication. Instead, they aimed at providing a model that ‘might be used in the design and implementation’ of local and temporary ‘patterns’ of interaction. That is, instead of deciding on a particular preconceived conception of CMC functionalities and applications, they ‘chose [...] to look at activities and the regulations required by a group of people to co-operatively execute a particular activity. The model we want to develop should therefore allow specification of such regulations’ (*ibid.*). That is, the aim was to build what one could call an abstract model or a notation that would make it possible ‘to model the activities, businesses, tasks, actions or workflow[s], which are performed by a group of co-operating people’, so as to, in turn, ‘facilitate the required co-ordination and possibly to automate co-ordination, thus reducing the co-ordination effort required of the participants in an activity’ (p. 23).

The European CMC researchers knew very well that the development of such computational models and architectures would have to be grounded in ‘fundamental understanding of Group Communication processes’ (p. 14), which in turn, because of the complexity and variability of working practices, would need contributions from ‘sociology, anthropology, economics and political science’ (p. 21). Their ‘strong commitment to the actual situation in working life’ was amply demonstrated in the pre-dominance of the practice-oriented program in the European CSCW research community that began to coalesce as these research activities ended in 1988. It is significant that Greif had reached strikingly similar conclusions: ‘Methodologies for testing individual user interfaces don’t apply as well to group support systems. As a result, CSCW is looking more to anthropology to find methodologies for studying groups at work in their natural settings’ (p.10).

In short, it was becoming clear that the CMC program was deeply flawed in its underlying ‘message handling’ outlook, in its focus on communication in the abstract, divorced from the work practices of which it normally is an integral part, but also severely limited in the way CMC conceived of the role of empirical studies in technological development. It was becoming clear, at least to some, that in-

depth studies of cooperative work practices in ‘natural settings’ was a prerequisite.

1.2. The rise and fall of Office Automation

At the same time as it was becoming clear to many CMC researchers, especially in Europe, that the ‘message handling’ paradigm was at odds with typical everyday cooperative work practices and that the paradigm thus had to be overcome, researchers in the ‘office automation’ movement were arriving at similar conclusions, although their point of departure was of course entirely different.

The ‘office automation’ movement had begun in high spirits in the 1970s, stimulated by different but intersecting technical developments. As with CMC, the baseline was the advent of computer networks. But the approach was radically different. Instead of conceiving of computer networks as a ‘medium’, that is, as a facility that regulates human interaction in negligible ways, the OA program deliberately aimed at regulating interaction in significant ways. The seminal idea was that various new techniques for constructing executable models that had been invented made it worthwhile to explore whether and to which extent such representations might be exploited as a means of modeling and regulating ‘office procedures’ and other kinds of workflows: on one hand, the algebraic techniques for building computational models of distributed systems developed by Petri and others since the early 1960s (cf., e.g., Zisman, 1977; Ellis, 1979) and, on the other hand, the equally sophisticated techniques for constructing complex adaptive models developed under the Artificial Intelligence label (cf., e.g., Hewitt, 1977; Fikes and Henderson, 1980; Barber and Hewitt, 1982). These hopes were soon defeated, however. Experimental applications such as DOMINO turned out to be felt like ‘straitjackets’ in actual use (Kreifelts, 1984; Kreifelts, *et al.*, 1991). Comparable lessons were learned from the CHAOS experiment (De Cindio, *et al.*, 1988; Bignoli and Simone, 1989). That is, ‘office work’ was not at all as easily captured and modelled as had been presumed. Handling contingencies and dealing with inconsistencies turned out to be an essential aspect of cooperative work practices. The ‘office automation’ program had landed in a crisis of its own.

At this point a new approach to technological research was devised: a few sociologists became involved in the effort to understand the status of ‘office procedures’ and cooperative work in general, on one hand Lucy Suchman and Eleanor Wynn (Wynn, 1979; Suchman, 1982, 1983; Suchman and Wynn, 1984) and on the other Eli Gerson and Susan Leigh Star (Gerson and Star, 1986).

That this coupling of sociological and technological research would first occur in the ‘office automation’ movement was hardly accidental. Email and most other CMC technologies were devised by computer technicians *for their own use*. That is, they were developed in a bottom-up and incremental fashion to solve local problems in practices that were well-known to the designers; and as they were found to be of general utility they were then — *post festum* — subjected to stan-

standardization and design. Their development did not require workplace studies of any kind. On the other hand, computer-conferencing systems were developed in a proactive manner; they were strictly speaking *designed*. But their design was based on normative models of what was claimed to be rational decision making, not on what was taken to be a well-grounded understanding of an actual practice. By contrast, however unrealistic the experimental designs of the ‘office automation’ movement turned out to be, nobody were under the illusion that one workflow model would fit all, and each workflow model was presumed to be empirically valid. That is, building technical systems that regulate actions and interactions in the strong sense envisioned by the ‘office automation’ movement was unproblematically thought to require some kind of analysis and modelling of existing procedures. When the models ultimately turned out not to work as anticipated, the natural next step was to look more carefully at the reality of ‘office work’.

This is anyway what happened. And it was also realized, eventually, that the problem was not just with this or that particular model or modelling technique. It was realized that the problem was conceptual. Those early studies of ‘office work’ indicated that received concepts of cooperative work as mere ‘execution’ of pre-conceived ‘procedures’ were inherently problematic. This point was driven home, emphatically, both by Gerson and Star and by Suchman in her contemporaneous critique of the concept of ‘plans’ in cognitive science (Suchman, 1987).

This insight was a fatal blow to the conceptual basis of the ‘office automation’ movement.

2. Enter CSCW

The work of Suchman, Wynn, Gerson and Star had significance beyond these, as it were, *immediate* implications. It also showed, *by way of example*, that not only were in-depth studies of actual working practices possible and fruitful; they also demonstrated that such studies could have strong impact on conceptual issues in technological research.

This, in my view, was the defining moment of CSCW. The early contributions by Wynn, Suchman, Gerson, and Star provided the ‘exemplars’, in a Kuhnian sense, for defining a new research program in which in-depth studies of cooperative work ‘in the wild’ were considered a prerequisite for developing computer technologies for human interaction. However, we should remember that new research paradigms are not necessarily heralded as such when they arrive on the scene. In fact, as pointed out by Kuhn, ‘we must recognize how very limited in both scope and precision a paradigm can be at the time of its first appearance’. Thus the ‘success of a paradigm [...] is at the start largely a promise of success discoverable in selected and still incomplete examples.’ (Kuhn, 1962, pp. 23 f.). This observation certainly applies to the emergence of the practice-oriented research program of CSCW.

The exemplary role of these studies were not only a function of the findings or of the role of field work in producing them. In both cases the research was integral to settings in which computer scientists and sociologists were addressing the same set of problems. The work of Suchman and Wynn was, of course, an important part of the research at Xerox PARC (from where the computer workstation and the Ethernet originated) where Suchman would later head a highly influential interdisciplinary group of researchers. It is less well known but important to note that the work of Gerson and Star anticipated much of what was later to unfold in CSCW in that their research was part of a collaborative research network involving both sociologists and computer scientists. The network, which *inter alia* also included Carl Hewitt, Anselm Strauss, Rob Kling, and Les Gasser, brought together sociologists with a track record in workplace studies of health care and biological research work *as well as* computer scientists engaged in developing what would later be known as distributed AI and agent-based architectures.

So, when Liam Bannon and I wrote our programmatic article for the first European CSCW conference in 1989, *this* was the kind of work we had in mind: ‘CSCW should be conceived as an endeavor to understand the nature and characteristics of cooperative work with the objective of designing adequate computer-based technologies. [...] The focus is to *understand*, so as to *better support*, cooperative work.’ (Bannon and Schmidt, 1989, p. 360).

In sum, two intellectual movements merged in the formation of CSCW. On one hand, CMC (as a technologically oriented research program) had arrived at a stage where it was beginning to dawn on many participants that the program was barking up the wrong tree. It had been focusing on aspects of interaction (‘communication’) that were conceived of as divorced from work practices but which normally are an integral part of doing the work and deeply enmeshed in the materiality of the setting and its organizational conventions and procedures. To move beyond that impasse, it was found necessary to develop an understanding of actual cooperative work practices. On the other hand, the ‘office automation’ program had landed in a situation where it had become clear that formal organizational constructs such as prescribed procedures are not mere algorithmic subroutines but part and parcel of professional work practices. It was, again, found necessary to develop an understanding of actual cooperative work practices. Here the history of CSCW proper begins.

When I point to the early work of Suchman, Wynn, Gerson, and Star as ‘exemplars’ of practice-oriented contributions to technological research, this of course does not mean that the formation of CSCW was not part of a wider intellectual movement than circumscribed by Ethnomethodology and Symbolic Interactionism. To the contrary. It was, and is, a distinct research effort within a much broader movement that, in different ways, strives to understand computing in its social context. Suffice it to mention the Participatory Design movement (e.g., Bjerknæs, *et al.*, 1987) that brought together computer scientists and others striv-

ing to understand the design and use of computing systems as embodied social practices. Likewise, subversive elements within Artificial Intelligence such as Terry Winograd quite early had serious doubts as to the conceptual foundations of AI and defected. At about the same time, a related movement away from cognitive science towards an ‘ecological’ and ‘naturalistic’ conception of computing was unfolding in Human Factors engineering. Consequently, when CSCW emerged as a distinct research program, it became a forum — and a rather tumultuous one at that — of these and other intellectual currents (Activity Theory, Distributed Cognition, etc.). When I nonetheless point to these early ‘exemplars’ it is because they, in different ways and from different intellectual traditions, demonstrated that in-depth studies of work practices could contribute not only to systems design but to the conceptual foundations of technological development.

1.1. CSCW’s program

What was new in CSCW, then? Firstly, the idea of doing field work as part of ‘requirements analysis’ is not at all new. The design of the very first computer applications for commercial purposes (payroll systems, etc.) was based on studies of actual practices. As early as 1953, the requirements analysis for one of the very first business applications, the design of a program for the ordering of goods for Lyons Teashops in the UK, involved genuine field work (Ferry, 2003, pp. 121-129). What was new in CSCW has to do with the difference between the *development of technologies*, i.e., technical knowledge, methods, principles, etc., and *systems design*, i.e., the instantiation and configuration of a set of known technologies for a specific purpose. The novelty of CSCW was not the idea of doing requirements analysis as an integrated part of the process of building a particular system for a particular setting, incorporating an array of more or less well-known technologies, but doing workplace studies for the purpose of developing new technologies, that is, to make field work an integral part of the conceptual work that is essential to technological research. Hence it is also not reasonable to expect of each and every study of particular work practices that it concludes in ‘implications for design’. What is required, however, is that studies should have *implications for CSCW*. The road from studies of work practices to technological development is indirect and complex. The role of ethnographic and other workplace studies in CSCW is not that of producing a requirements analysis but to contribute conceptually.

The fecundity of CSCW’s practice-oriented program became evident immediately, even as the program was being tentatively articulated. The first report on the Lancaster group’s study of air traffic control was presented to the incipient CSCW community in 1989 (Harper, *et al.*, 1989) and was quickly followed by the equally emblematic study of the London Underground control room (Heath and Luff, 1991). Nor did it not take long for it to become clear that these new insights would have radical implications for not only the development of certain classes of

applications but for underlying computer technologies. This was, for example, made explicit with respect to the research area of distributed systems by Rodden and Blair in their classic paper from 1991. Referring to the ‘the rich patterns of cooperation found in CSCW’ depicted in the early harvest of ethnographic studies in CSCW, the authors summarized their programmatic argumentation by stating that ‘existing approaches to control in distributed systems are inadequate’ (Rodden and Blair, 1991, p. 49). The implications for technological research are profound:

‘For example, consider the problem of shared access to resources. In most distributed systems this is dealt with by masking out the existence of other users. Hence sharing is transparent with each user unaware of the activity of others. This clearly contradicts the needs of CSCW. [...] The problem with this approach is that presumed control decisions are embedded into the system and hence cannot be avoided or tailored for specific classes of application. This is the root of the problem in supporting CSCW. Because of the dynamic requirements of CSCW applications, it is very unlikely that such prescribed solutions will be suitable.’ (Rodden and Blair, 1991, p. 59)

Rodden and Blair concluded that ‘CSCW demands a fresh approach to control which is specifically tailored for cooperative working’ (Rodden and Blair, 1991, p. 60). This was a crucial programmatic proposition. The key problem for CSCW is not ‘communication’ or ‘resource sharing’ but the cooperating actors’ control of their interaction and, by implication, of the computational regulation of their interaction. This problem is fundamentally different from the issue of user control of system behavior in HCI, in that control in cooperative work settings is, in principle, distributed. This problem has since then been spelled out and elaborated under from different perspectives: ‘event propagation mechanisms’ for ‘awareness’ support, ‘coordinative artifacts and protocols’, and so on.

As observed above, the paradigm of the research program that is CSCW was exemplified by the early studies by Wynn, Suchman, Gerson, and Star which demonstrated how sociological inquiries could address conceptual issues in technological research. Similarly, with Rodden and Blair’s re-conceptualization of fundamental issues in distributed computing CSCW’s research program had been complemented by an exemplar of the correlative technological research. The reciprocity of the contributions of sociology and computer science respectively had also been exemplified.

1.2. The afterlife of CMC in CSCW

The developments within CMC research after the crisis and the formation of the CSCW research program are complex.

In a sense, the pattern of original CMC technology development has repeated itself, again and again. The case in point is of course the development of the Web (HTTP and HTML). It was initially developed by scientists at CERN for their own use, and the initial motive was almost identical with that of the Internet: ac-

cess to resources across platforms. The technologies were themselves derived from previous technologies such as hypertext and markup languages. However, as with network email, when it arrived the Web was soon adopted by others to be used in other contexts (Gillies and Cailliau, 2000).

The pattern is characterized by occasional technological innovation, innovative applications of well-known technologies, often in novel configurations, and a significant element of reimplementation for other purposes in new contexts. As a result, wave upon wave of seemingly new communication facilities have, again and again, caught the attention of the media and the public at large: instant messaging, text messaging, chat, blogs, and so on. Some of them, such as instant messaging go back to the early days of time-sharing operating systems. What is new, however, is that they have been somewhat standardized so that they can be used across different platforms and, consequently, have been adopted by a mass audience. Similarly, in the case of chat and blogs, we are talking about facilities that are re-implementations of computer-conferencing and 'bulletin boards' *anno* 1980. What has given the scaled-down computer-conferencing idea a new lease on life is the ubiquity of the web browser: the HTTP protocol has become a general platform-independent way of establishing conversational sessions that are then governed by other communication protocols. And again, as a result of the 'super-platform' provided by the web browser, these conferencing facilities have been picked up on a mass scale too. That is, what is generally happening is that well-known computer-based communication technologies, often in innovative configurations, are reaching a mass audience.

The relationship of CMC research to these developments is not less complicated. But typically the new implementations have been undertaken by designers for their own use in their own particular part of the woods, only to be picked up and spread in a classical innovation-diffusion pattern. In a way that is reminiscent of the pattern of the '70s and early '80s, CMC research, in its many forms, strives to investigate possible 'effects' and 'impacts' of these socio-technical phenomena, but without the original's relatively close coupling to experimental technological development. This is not surprising, since the socio-technical phenomena under investigation typically do not represent new technologies but rather new applications on a mass scale.

It is here important to point out that, over the last ten or fifteen years, a large and diverse area of research, normally *also* referred to as CMC research, has emerged that does not have any relationship with technological research and does not consider itself related to the concerns and issues of CSCW. To use the wording of the program statement of one of the leading journals in this field, this research 'is concerned with the empirical study of human behavior in the online environment, and with the impact of evolving communication and information tech-

nology upon individuals, groups, organizations, and society'.¹ — The reason why this research area refers to itself as CMC research is simply that CMC technologies underpin the media that facilitate the behavior under investigation, just like other fields of communication and media studies investigate behavior connected to movies or TV. Such inquiries may be worthwhile, although inquiries that aim at understanding or even anticipating the societal impact of technical innovations are in a methodological muddle: 'much of the CMC work still holds to an overly determinist view concerning the role of technology in human affairs, attempting to assess impacts of new technology in general, and missing out on the interplay of social forces in the acceptance and use of the new media' (Bannon, 1992). Anyway, this body of CMC research addresses problems and conceives of its findings in ways that have no direct bearing on CSCW, and it is not my concern here.

My concern here is with the fact that CMC research continues within CSCW unaware of and unaffected by the fact that its conceptual legitimacy has been fundamentally challenged. Although the CMC paradigm's focus on communication as a separate activity has been found wanting and its model 'broken', and although this realization, together with the experience of the OA program, has given rise to a new research program with an entirely different paradigm, CMC research has continued in CSCW unabated, as if nothing has happened. In fact, it is becoming predominant, at least in quantitative terms.

Now, in so far as CMC facilities are adopted in work settings, which they obviously are, they may of course be of some interest to CSCW, as facilities we can build on or otherwise have to relate to — on par with database systems, network facilities, modelling techniques, or sensor and actuator technologies. Furthermore, the appropriation and use of CMC facilities and techniques in work settings may raise many interesting issues. These facilities are, for example, being deployed in ways that may change organizational boundaries and roles, blur the traditional separation of work and leisure in terms of time and place, and so on. These are issues that occupy researchers in economics, occupational sociology, and organization theory but they do have implications for CSCW in as much as they affect the organizational and material settings of cooperative work.

It is, in this context, also of relevance that net-based communication facilities are being employed to enable increasing geographical distribution of work in the form of, e.g., global production networks. These developments raise questions concerning the organization and management of cooperative work in such dispersed settings (cf. Hinds and Kiesler, 2002). These are important issues. Indeed, coordinating interdependent activities across space is one of the problems faced by actors engaged in cooperative work 'in the wild'. However, the model of com-

¹ *The Journal of Online Behavior*: 'Overview' (<http://www.behavior.net/JOB/job.html>). According to the journal's editors, topics typically investigated in this area are: 'The role of the Internet in national and local news media use', 'The relationship between exposure to Internet pornography and sexual attitudes toward women', and 'Reformulating the Internet paradox: Social cognitive explanations of Internet use and depression'.

puter-mediated communication ‘breaks down’ (again) when the issue is investigated in abstraction from the actual coordinative practices of, say, software engineering work. And at any rate, handling cooperative interaction across geographical distance is only one issue in the coordination of interdependent activities, and it is thus absurd to define CSCW in terms of the issue of distance.

That is, for studies of CMC facilities to contribute to the technological commitments of CSCW, they would have to investigate how these facilities are appropriated in actual coordinative practices, that is, how practitioners integrate these facilities with their repertoire of coordinative artifacts, in their embodied activities, in material work settings. However, this kind of investigation falls well outside of what could be called ‘the divorced-communication paradigm’ that characterizes CMC research in CSCW.

What characterized CMC research in CSCW is, first of all, that it conceives of communication in abstraction from actual cooperative work practices. This is a fundamental precept inherited from the original CMC research. It defines its ‘world view’: what is considered relevant and perhaps even researchable. Secondly, CMC research focuses on computer networks as a means that facilitates interaction with only rudimentary computational regulation, as facilities on par with television and radio or telegraph and telephone (Olson and Olson, 2003, p. 584). In that sense, the program is faithful to the received ‘medium’ metaphor. But at that level of abstraction, e.g., in terms of ‘media characteristics’, no contribution to the development of technology is possible. Thirdly, however, in contrast to original CMC research, CMC in CSCW is reactive, conceiving of empirical work as something *post hoc* to technological development (as a kind of technology assessment).

As described above, CMC research formed in the ’70s to investigate the new communication technologies that were being developed by computer technicians for their own use (in building, operating, maintaining whatever software and hardware systems they were working on) or in some cases deliberately designed for the use by others. CMC research anyway formed in close coupling with these development activities, sometimes carried out by the technicians or at least in close collaboration with the technicians. However, as email and the other forms of CMC technologies became standardized services, CMC research was left dangling. But when new applications of CMC technologies began to occur, especially spurred by the emergence of the Web, the methods and techniques of the original CMC research program were found applicable again, only now the continuation of CMC evaluation work had lost its connection to design and became a special blend of technology assessment and technology transfer. As pointed out by Bannon in 1992, ‘the orientation of much of the CMC work is on evaluation rather than on gathering material to be used for design or re-design of technologies. Research is thus more reactive than pro-active. This affects the kinds of research methods used, the problems addressed in research, etc.’ (Bannon, 1992).

What has remained constant in CMC research is the ‘the divorced-communication paradigm’. Thus, for the purposes of CSCW, CMC research is marginal at best, a distraction at worst. In sum, CSCW as a community comprises not only different research programs but incommensurate paradigms.

What, then, drives the CMC research program in CSCW? An obvious reason is of course the wave after wave of seemingly new CMC facilities that fascinate the public and researchers alike. Another, perhaps supplementary, explanation of the unabated continuation of CMC evaluation studies in CSCW would be that many researchers have retained strong disciplinary reservations towards ethnography and other forms of workplace studies. For example, some CSCW researchers claim that the central role that ethnographic studies of actual practices holds in CSCW is in fact a source of ‘weakness’ of CSCW, and they advocate a ‘stronger orientation’ to what is claimed to be ‘a large body of well-validated principles about human behavior in group and organizational contexts’ that, correspondingly, employs ‘data collection and analysis methods that emphasize parsimony and identification of generalizable features of human behavior’. The aim of this, they state, is to develop ‘universal principles of CSCW design’ (Finholt and Teasley, 1998, p. 40 f.).

A discussion of the assumptions underlying this criticism of the role of ethnography in CSCW is of course beyond the scope of this paper. But a couple of points need to be made. First of all, it would of course be absurd to claim that just because a particular study of a cooperative work settings is based on ‘ethnography’, for instance by virtue of somebody’s having been at the site for some time and observed events, then it is a valid contribution to CSCW, whereas a study that employs quantitative techniques is ruled out. Whatever the actual investigative technique, the issue is rather the specific analytical stance of the study: Does it provide an in-depth analysis of the *logic of the work practice* in question? Having said that, I should point out that Finholt and Teasley seem to take for granted, without reflection or argument, that there is one and only one legitimate form of scientific generalization, namely that of identifying abstract universal principles (e.g., ‘laws’). Such an assumption is not only evidently false, as it would outlaw scientific insights of great value in a range of research fields. But in our context such dogmatism would lead to impotence. Indeed, in the context of cooperative work practices, such abstractions would be meaningless. Let me be a little more specific. Their criticism of the dominant role of ethnography in CSCW begs the question how one, on the basis of ‘universal principles of CSCW design’, can devise technologies that regulate *historically specific* professional work practices or take account of what Rodden and Blair called the ‘rich pattern of cooperation’. Could general ‘principles about human behavior in group and organizational contexts’ tell us *anything relevant* about the contingent handling of production planning systems in manufacturing, the development of naming schemes in engineering, the coding practices of medical records, the role of flight progress strips in air

traffic control? But such questions are of course inconceivable within the CMC world view.

1.3. Implications for design...

George Bernard Shaw is often cited for having remarked that ‘Britain and America are two nations divided by a common language’. The same is true of CSCW. But here the confusion has implications of greater import than the occasional misunderstanding between Americans and Brits.

The fragmentation of CSCW is harmful. While mutual indifference between different schools may be acceptable in a field that does not aim at contributing to the development of technology, it is fatal to a field like CSCW. It fosters confusion and discontinuity; it makes it exceedingly difficult for the field to work in a cumulative or converging manner. The effect of that, in turn, is that CSCW is seriously handicapped in meeting its commitment to the development of computer-based technologies by means of which members of ordinary work settings can control the computational coordination of their distributed and yet interdependent activities.

Acknowledgements

The research has been supported by the Danish Research Council for Technology and Production (CosmoBiz, Grant no. 274-06-0415) and the Danish Strategic Research Council (CITH, Grant No. 2106-07-0017). I thank the six anonymous reviewers as well as the two ‘shepherds’ appointed by the Program Committee for criticisms and comments that helped me to make the thrust of the article come across more clearly. Last, but not least, I would like to express my gratitude to Liam Bannon for trying to instill in me a sense of the finer nuances of CMC research.

References

- Bannon, Liam J.: ‘Computer-Mediated Communication’, in D. A. Norman and S. W. Draper (eds.): *User Centered System Design*, Lawrence Erlbaum, New Jersey, 1986, pp. 433-452.
- Bannon, Liam J.; Niels Bjørn-Andersen; and Benedicte Due-Thomsen: ‘Computer Support for Cooperative Work: An appraisal and critique’, in *Eurinfo’88: 1st European Conf. on Information Technology for Organisational Systems, 16-20 May 1988, Athens, Greece*, 1988.
- Bannon, Liam J.; and Kjeld Schmidt: ‘CSCW: Four characters in search of a context’, in *ECSCW’89, 13-15 Sept. 1989, Gatwick, London*, London, 1989, pp. 358-372.
- Bannon, Liam J.: ‘Perspectives on CSCW: From HCI and CMC to CSCW’, in *EW-HCI’92: Proc. Int. Conf. on HCI, August 1992, St. Petersburg, Russia*, 1992, pp. 148-158.
- Barber, Gerald R.; and Carl Hewitt: ‘Foundations for office semantics’, in N. Naffah (ed.): *Office Information Systems*, INRIA/North-Holland, Amsterdam, 1982, pp. 363-382.
- Bignoli, Celsina; and Carla Simone: ‘AI Techniques for supporting human to human communication in CHAOS’, in *ECSCW’89, 13-15 Sept. 1989, Gatwick, London*, London, 1989, pp. 133-147.
- Bjerknes, Gro; Pelle Ehn; and Morten Kyng (eds.): *Computers and Democracy : A Scandinavian Challenge*, Avebury, Aldershot, 1987.

- Bowers, John M.; and Steven D. Benford: 'Part one: The concept of CSCW', in Bowers and Benford (eds.): *Studies in Computer Supported Cooperative Work: Theory, Practice and Design*, North-Holland, Amsterdam, 1991, pp. 1-2.
- Crocker, David H., et al.: 'RFC 733: Standard for the format of ARPA network text messages', 21 November 1977. <<ftp://ftp.rfc-editor.org/in-notes/rfc733.txt>>
- De Cindio, Fiorella, et al.: 'CHAOS: a knowledge-based system for conversing within offices', in W.Lamersdorf (ed.): *Office Knowledge: Representation, Management and Utilization*, Elsevier North-Holland, Amsterdam, 1988.
- Ellis, Clarence A.: 'Information control nets: A mathematical model of office information flow', in: *Proc. ACM Conf. on Simulation, Measurement and Modeling, August 1979, Boulder, Colorado*, ACM, New York, 1979, pp. 225-239.
- Ferry, Georgina: *A Computer Called LEO: Lyons Teashops and the World's First Office Computer*, HarperCollins, London, 2003. (Paperback ed., 2004).
- Fikes, Richard E.; and D. Austin Henderson, Jr.: 'On supporting the use of procedures in office work', in *Proc. 1st Ann. Conf. on AI, 18-20 August 1980, Stanford Univ., AAAI, 1980*, pp. 202-207.
- Finholt, Thomas A.; and Stephanie D. Teasley: 'Psychology: The need for psychology in research on computer-supported cooperative work', *Social Science Computer Review*, vol. 16, no. 1, Spring 1998, pp. 40-52.
- Gerson, Elihu M.; and Susan Leigh Star: 'Analyzing due process in the workplace', *ACM Transactions on Office Information Systems*, vol. 4, no. 3, July 1986, pp. 257-270.
- Gillies, James; and Robert Cailliau: *How the Web was Born: The Story of the World Wide Web*, Oxford Univ. Press, Oxford, 2000.
- Greif, Irene: 'Overview', in I. Greif (ed.): *Computer-Supported Cooperative Work: A Book of Readings*, Morgan Kaufmann Publishers, San Mateo, Calif., 1988, pp. 5-12.
- Hardy, Ian R.: *The Evolution of ARPANET email*, History thesis paper, Univ. of Calif., Berkeley, 13 May 1996. <<http://www.ifla.org.sg/documents/internet/hari1.txt>>
- Harper, Richard; John A. Hughes; and Dan Shapiro: 'Working in harmony: An examination of computer technology in air traffic control', in *ECSCW'89, 13-15 Sept. 1989, Gatwick, London*, London, 1989, pp. 73-86.
- Heath, Christian C.; and Paul Luff: 'Collaborative activity and technological design: Task coordination in London Underground control rooms', in *ECSCW'91, 24-27 Sept. 1991, Amsterdam*, Kluwer, Dordrecht, 1991, pp. 65-80.
- Hewitt, Carl: 'Viewing control structures as patterns of passing messages', *Artificial Intelligence*, vol. 8, 1977, pp. 323-364.
- Hiltz, Starr Roxanne; and Murray Turoff: 'The evolution of user behavior in a computerized conferencing system', *CACM*, vol. 24, no. 11, November 1981, pp. 739-751.
- Hiltz, Starr Roxanne; and Murray Turoff: 'Structuring computer-mediated communication systems to avoid information overload', *CACM*, vol. 28, no. 7, July 1985, pp. 680-689.
- Hinds, Pamela J.; and Sara Kiesler (eds.): *Distributed Work*, MIT Press, Cambridge, Mass., 2002.
- Kerr, Elaine B.; and Starr Roxanne Hiltz: *Computer-Mediated Communication Systems: Status and Evaluation*, Academic Press, Orlando, etc., 1982.
- Kreifelts, Thomas: 'DOMINO: Ein System zur Abwicklung arbeitsteiliger Vorgänge im Büro', *Angewandte Informatik*, vol. 26, no. 4, 1984, pp. 137-146.
- Kreifelts, Thomas, et al.: 'Experiences with the DOMINO office procedure system', in *ECSCW'91, 24-27 Sept. 1991, Amsterdam*, Kluwer, Dordrecht, 1991, pp. 117-130.
- Kuhn, Thomas S.: *The Structure of Scientific Revolutions* (1962; 2nd ed. 1969). Univ. of Chicago Press, Chicago, 1969.
- Lakatos, Imre: 'Falsification and the methodology of scientific research programmes', in I. Lakatos and A. Musgrave (eds.): *Criticism and the Growth of Knowledge*, Cambridge Univ. Press, Cambridge, 1970, pp. 91-196.

- Licklider, J. C. R.: 'Man-computer symbiosis' (*IRE Transactions on Human Factors in Electronics*, March 1960). In R. W. Taylor (ed.): *In Memoriam: J. C. R. Licklider, 1915-1990*. Digital Systems Research Center, Palo Alto, Calif., 1990, pp. 4-11.
- Licklider, J. C. R.; and Welden E. Clark: 'On-line man-computer communication', in *AFIPS Spring Joint Computer Conf.*, vol. 21, 1962, pp. 113-128.
- McCarthy, John: 'Reminiscences on the history of time sharing', Stanford Univ., Winter or Spring 1983. <<http://www-formal.stanford.edu/jmc/history/timesharing/timesharing.html>>
- Myer, Theodore H.; and David Dodds: 'Notes on the development of message technology', in *Berkeley Workshop on Distributed Data Management and Computer Networks, 1976*, Lawrence Berkeley Laboratories, 1976, pp. 144-154. – LBL-5315.
- Olson, Gary M.; and Judith S. Olson: 'Groupware and computer-supported cooperative work', in J. A. Jacko and A. Sears (eds.): *The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, Lawrence Erlbaum, Mahwah, New Jersey, 2003, pp. 583-595.
- Palme, Jacob: 'You have 134 unread mail! Do you want to read them now?', in *IFIP Conf. on Computer Based Message Services, Nottingham*, Elsevier, New York, 1984, pp. 175-184.
- Panko, Raymond R.: 'The outlook for computer mail', *Telecommunications Policy*, June 1977, pp. 242-253.
- Pankoke-Babatz, Uta (ed.): *Computer Based Group Communication: The AMIGO Activity Model*, Ellis Horwood Publishers, Chichester, 1989.
- Postel, Jon: 'RFC 808: Summary of computer mail services meeting held at BBN on 10 January 1979', 1 March 1982. <<ftp://ftp.rfc-editor.org/in-notes/rfc808.txt>>
- Rodden, Tom A.; and Gordon Blair: 'CSCW and distributed systems: The problem of control', in *ECSCW'91, 24-27 Sept. 1991, Amsterdam*, Kluwer, Dordrecht, 1991, pp. 49-64.
- Suchman, Lucy A.: 'Systematics of office work: Office studies for knowledge-based systems. Digest', in *Office Automation Conf., 5-7 April 1982, San Francisco*, 1982, pp. 409-412.
- Suchman, Lucy A.: 'Office procedure as practical action: Models of work and system design', *ACM Transactions on Office Information Systems*, vol. 1, no. 4, October 1983, pp. 320-328.
- Suchman, Lucy A.; and Eleanor Wynn: 'Procedures and problems in the office', *Office: Technology and People*, vol. 2, 1984, pp. 133-154.
- Suchman, Lucy A.: *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge Univ. Press, Cambridge, 1987.
- Turoff, Murray: 'Delphi conferencing: Computer-based conferencing with anonymity', *Technological Forecasting and Social Change*, vol. 3, 1972, pp. 159-204.
- Turoff, Murray: 'Human communication via data networks', *SIGCAS Computers and Society*, vol. 4, no. 1, May 1973, pp. 15-24.
- Vallee, Jacques F.: 'The FORUM project: Network conferencing and its future applications', *Computer Networks*, vol. 1, 1976, pp. 39-52.
- Van Vleck, Tom: 'The history of electronic mail', 1 February 2001. <<http://www.multicians.org/thvv/mail-history.html>>
- Wynn, Eleanor H.: *Office Conversation as an Information Medium*, Ph.D. diss., Univ. of California, Berkeley, 1979.
- Zisman, Michael D.: *Representation, Specification and Automation of Office Procedures*, Ph.D. diss., Univ. of Pennsylvania, 1977.

Collaboration in Metagenomics: Sequence Databases and the Organization of Scientific Work

Matthew J. Bietz & Charlotte P. Lee

University of Washington

{mbietz, cplee}@u.washington.edu

Abstract. In this paper we conduct an ethnographic study of work to explore the interaction between scientific collaboration and computing technologies in the emerging science of metagenomics. In particular, we explore how databases serve to organize scientific collaboration. We find databases existing across scientific communities where scientists have different practices and priorities. We suggest while these databases appear to be boundary objects, they are better understood as boundary negotiating artifacts. Due to rapid scientific and technical innovation the tools, practices, and scientific questions change over the course of merely a few years resulting in challenges for collaboration.

Introduction

The use of databases is critical for metagenomic science. While databases are often intended to span the boundaries between communities of practice (Wenger, 1998), they actually serve more as sites for the negotiation of scientific methods, research questions, and worldviews. Due to rapid scientific and technical innovation the tools, practices, and scientific questions change over the course of merely a few years. We find that multiple databases are useful for supporting work in a highly dynamic context of leading edge science. It is the production and use of these databases, and the implications therein for collaborative work and technology design, that will concern us in this paper.

Unlike traditional genomics, which focuses on the genetic code of specific organisms or species, metagenomics focuses on the distribution of genetic material within a population of microorganisms. The move toward metagenomic approaches has been enabled by technological advances in DNA sequencing that have allowed the generation of large amounts of data at significantly lower cost. These new technologies have created a number of cyberinfrastructure-related challenges, including exponentially increasing computational power and data storage needs and new algorithms for manipulating and analyzing data. Metagenomics also requires an interdisciplinary approach, frequently bringing together ecologists, geneticists, bioinformaticists, and computer scientists.

Our research suggests that metagenomics researchers, bioinformaticists, and developers of cyberinfrastructure often have a strong sense of an ideal conceptual Database that would contain all genetic sequence data and associated metadata, and often derivative analyses. This conceptual Database and implemented database systems frequently serve as a boundary negotiating artifacts (Bowker & Star, 1999; Lee, 2007). The multiplicity of databases, in particular, play a useful role in supporting highly innovative and dynamically changing activities.

Background: Cyberinfrastructure and Databases

Cyberinfrastructures are distributed organizations supported by advanced technological infrastructures such as supercomputers and high-speed networks. Within the domain of scientific cyberinfrastructures (also known as e-Science), the capabilities of cyberinfrastructure are usually intended to be transformative (Atkins, et al., 2003). Cyberinfrastructures are employed to support work, often in the form of collaborative data sharing and, less frequently, analysis. The ability to pool data can enable scientists to answer questions that no single investigator or laboratory could answer individually. Large-scale data sharing can enable not only new types of scientific practices, but can also enable the exploration of new types of scientific questions. Conducting new types of science requires new and more powerful technologies to support communication, data sharing and analysis, and access to remote instruments.

The sharing of data, however, is rarely straightforward. Previous research has shown that the development of effective CSCW systems to support data sharing groups requires a better understanding of the use of data in practice. Data play two general roles in scientific communities: 1) they serve as evidence to support scientific inquiry, and 2) they make a social contribution to the establishment and maintenance of communities of practice (Birnholtz & Bietz, 2003). Birnholtz and Bietz found that data sharing, particularly in fields with high task uncertainty, is problematic because of the difficulty of communicating contextual information in the absence of interpersonal interaction. Needed

contextual information includes the nature of the data, the scientific purpose of its collection, and the social function in the community that created it.

Issues of data sharing are critical to the development of large scale information infrastructures, but a treatment of data sharing should also engage a discussion of databases. In her work on databases as scientific instruments, Hine (2006) found a mouse genome database to be an emergent structure that is necessarily embedded in particular sets of work practices. She notes that:

The patterns of connection and collaboration in scientific knowledge production involving databases can thus become both spatially and socially complex, building on existing networks but adding additional density, bandwidth and new tensions (Hine, 2006, p. 293).

In other words, the database is both built upon and enabling of scientific collaboration. The database provides both opportunities and constraints.

Boundary Negotiating Artifacts

CSCW has long studied coordinative artifacts for the purposes of theorizing collaboration as well as informing the design of groupware. Many types of artifacts have coordinative functions and databases (not just the information contained therein) may be included among these. Research on coordinative artifacts have focused on paper and electronic documents (Lutters & Ackerman, 2002; Schmidt & Simone, 1996; Schmidt & Wagner, 2002) and have looked at these documents as boundary objects (Bowker, et al., 1999; Star, 1987-1989), boundary negotiating artifacts (Lee, 2007), and have put forth useful methodologies with which to understand how documents such as a report can serve to organize work in the most complex of organizations (Harper, 1998; Schmidt & Wagner, 2005). Many of these papers include in their purview spreadsheet documents and databased information (Harper, et al., 2001; Lee, 2007). Other fields have emphasized that the database itself is an important cultural form that entails a different mode of thinking about the world (Manovich, 2001) and as occasioning a new set of arrangements, as opposed to scientific journals for example, for the communication of scientific information and methods (Hilgartner, 1995).

Previous research has defined a *shared information system*, such as a shared database, to be an information system that is used by multiple communities of practice (Pawlowski, et al., 2000). These systems are described as typically spanning formal organizational boundaries such as functional departments or business units. Pawlowski et. al (2000) focus on enterprise-wide databases that support beginning-to-end business processes. They note that maintaining a shared system in an organization is challenging because triggers for system change can originate in any of the stakeholder areas when work practices or requirements change. Ultimately they suggest that shared information databases and related artifacts are boundary objects that require brokering, translating, coordinating and aligning perspectives, and addressing conflicting systems. Although we agree with

the larger premise that databases that are used by multiple communities of practice are key for boundary work and that these databases require brokering, a careful reading of the original work on boundary objects (Star, 1987-1989; Star & Griesemer, 1989) suggest that these databases may actually be a combination of boundary negotiating artifacts and boundary objects, or they may simply be boundary negotiating artifacts.

Defining features of boundary objects include that they pass from one community of practice to another with little or no explanation and satisfy the informational requirements of multiple communities of practice. Yet some of the things we call boundary objects do not actually do so (Lee, 2007). Throughout the literature described above, the following themes recur: So-called boundary objects may require considerable additional explanation and discussion to be intelligible; Artifacts sometimes play a role in the *active negotiation* of shared understanding amongst communities of practice (and thus can be used to enlist participation and can be adjusted through group interaction); Unstandardized artifacts that are partial, incomplete, or are intermediary representations are ubiquitous in collaborative work; And so-called boundary objects can “fail” to satisfy the informational needs of collaborating parties (Henderson, 1999; Lee, 2007; Subrahmanian, et al., 2003). The recurring themes described here indicate that the boundary objects concept is not incorrect, rather it is incomplete. Other researchers have grappled with fitting their research findings to the notion of boundary objects. Henderson (1999) found that the boundary object concept required amendment and suggested the term *conscription devices* to refer to a type of boundary object that enlists group participation and that are adjusted through group interaction. Subrahmanian et al. (2003) proposed the broad concept of *prototypes* based on their observation of artifacts that support systematic updating of boundary objects. Organizational changes, they note, sometimes caused boundary objects to be unable to support activity. O’Day et al. (2001), in their work on molecular biologists and computer scientists, refer to *boundary objects in-the-making* which are unstable objects that still work to facilitate collaboration across communities by giving people common ground for discussion and negotiation. They note that in the absence of durable cooperation, boundary objects in-the-making are necessary to confront and reconcile different local meanings.

We stress the importance of adopting a strict definition of boundary objects that is true its origins. By doing so, we can fully appreciate just how large and nuanced is the research and design space when we accept the idea that many artifacts and practices are not just crossing but weaving, pushing, pulling, and everything else on, around, and through communities of practice. Boundary negotiating artifacts provide a lens through which we can view the myriad uses of artifacts, many of them messy and ad hoc. Boundary negotiating artifacts:

- Are surrounded by sets of practices that may or may not be agreed upon by participants

- Facilitate the crossing of boundaries (transmitting information)
- Facilitate the pushing and establishing of boundaries (dividing labor)
- May seem “effortful” in use as opposed to effortless
- Are fluid: 1) a boundary negotiating artifact can change from one type to another when the context of use changes; and 2) a boundary negotiating artifact can sometimes also simultaneously be physically incorporated or transformed into another artifact
- Can be largely sufficient for collaboration
- Are possible predecessors of boundary objects (Lee, 2007)

Boundary negotiating artifacts are used to: record, organize, explore and share ideas; introduce concepts and techniques; create alliances; create a venue for the exchange of information; augment brokering activities; and create shared understanding about specific problems. Scientific collaboration between biologists and computational disciplines have been noted as an endeavor that requires interpretive frames to be brought together:

At the end of the day, people in biological and computational disciplines try to produce biological understanding by bringing their distinctive interpretive frames together. But as we have discussed, it is likely that there will be an ongoing need for negotiation between disciplines. It is not the case that biologists can simply learn how to run the numbers; the numbers and way to run them continue to be problematic as biologists ask new questions and encounter new forms of data. Similarly, mathematicians and computer scientists are challenged to develop new analytical methods to deal with the flexibility and multi-dimensionality of living systems. Biologists and computation experts need to continue their collaboration. (O'Day, et al., 2001, p. 417)

In the research that follows, we too find that the practices and interpretive frames of biologists and computational disciplines need to be brought together, but in the area of metagenomics research there is even greater collaborative complexity. We find that the communities of biologists who are attempting to share databases have diverse practices, interpretive frames, and different scientific concerns that are brought together by databases which function as boundary negotiating artifacts.

The Current Study

We employed ethnographic research methods which involved entering into sites involved in the production of metagenomics research and databases, getting to know the people involved, participating in the daily routines of the setting, and observing what is going on. Our goal was to observe ordinary conditions, responses to events, and experience events ourselves as much as possible in order to understand “social life as process” (Emerson, et al., 1995).

Our engagement with these communities began in the summer of 2007, and is ongoing at the time of writing. Our initial focus was on one particular metagenomics database project, and our primary focus for the first year of our engagement was on the team developing the database. We interviewed as many members of the development team as we could, some of them multiple times. For four months of this time, one of the authors attended weekly project meetings, ad

hoc meetings, and spent at least one day per week working from an assigned desk in the development team area.

In the second year of engagement, our focus shifted to developing a broader understanding of the landscape of cyberinfrastructure for metagenomics research. We interviewed microbiologists, bioinformaticists, computer scientists, and representatives of funding agencies. We interviewed both users and developers of several major genomics and metagenomics databases. We attended conferences and workshops devoted to metagenomics research, database development, and the development of standards. For the past seven months, one of the authors has attended weekly laboratory meetings at an academic molecular biology laboratory engaged in metagenomics research.

In total, this amounts to thirty-three formal interviews and well over one-hundred hours of on-site observation and informal conversation. Interviews were semi-structured and ranged from thirty minutes to nearly two hours, with most lasting between sixty and ninety minutes. Transcriptions of the interviews, field notes, and various indigenous documents were coded in Atlas.ti using a grounded theory approach (Glaser & Strauss, 1967).

A Metagenomics Primer

The term “metagenomics” was coined in 1998, and while there is some controversy about the exact definition of the term, it generally refers to using genomics techniques to study communities of microorganisms (Chen & Pachter, 2005; Handelsman, et al., 1998). Until recently, it was necessary to culture microorganisms in a laboratory in order to produce enough DNA for sequencing. However, it is estimated that less than one percent of the world’s microorganisms can be cultured using standard laboratory techniques (Hugenholtz, et al., 1998). Advances in DNA amplification techniques and new sequencing technologies have significantly reduced the cost of sequencing and made it possible to analyze DNA without culturing, giving scientists access to a newfound wealth of genetic information.

Metagenomic techniques are relevant to a number of fields, including marine ecology, medicine, energy production, and environmental remediation, to name a few. In a typical metagenomic experiment, scientists begin by sampling the microorganisms from a particular environment. For example, a marine microbiologist may pass seawater through a series of progressively smaller filters to isolate a particular kind of microorganism (viruses, bacteria, etc.). DNA is then extracted from the organisms and prepared for sequencing. While the specifics vary across manufacturers and technologies, many metagenomic analyses use “shotgun sequencing,” in which the long strands of DNA are randomly broken up into shorter segments which are “read” by the sequencer. Depending on the technology, these “reads” range from 20 to 400 base pairs in length. Longer

segments of DNA are then computationally reconstructed by searching for areas of overlap among the shorter segments.

Because of the number of different organisms in the sample, this technique typically results in only a small portion of each organism's genome being sequenced. Through a combination of statistical techniques and comparisons to known genomes, scientists can identify the most prevalent organisms in their sample and estimate the diversity of organisms in the population. Scientists can also study the functional capacities of the population of microorganisms in relation to their environment, for example, the ability of marine microbes to metabolize phosphorous (Gilbert, et al., 2009) or the influence of the microbial population of the gut on obesity (Turnbaugh, et al., 2008).

Computation in Metagenomics

A discussion of metagenomics practice would be incomplete without a discussion of the computational resources on which metagenomics relies. Metagenomics would not be possible without a broad array of computational tools and information systems. Computation is so central to the work of these scientists that most of our senior biologists spent very little time at “the bench” working with wet materials. Although their students spent more time at the bench and in the field collecting samples, our senior participants all reported spending at least 90% of their research time at a computer.

Our participants often resort to the metaphor of jigsaw puzzles to explain the role of computation in metagenomics. Environmental shotgun sequencing has been compared to mixing the pieces from many different jigsaw puzzles in the same bag and pulling out a few handfuls of the pieces. The computer is used to put the pieces together when possible, and from the resulting fragments of puzzles, try to figure out how many puzzles were in the bag and what picture was on each one.

So the first computational task is to assemble the fragmented DNA sequences (the puzzle pieces) into longer contiguous sequences. This is made more complicated by differences among sequencing technologies, which result in varying read lengths and error rates. Even so, assembly is seen to be a relatively straightforward process compared to the later analysis of the assembled sequences.

Frequently the next step is to make the assembled sequences biologically meaningful by “annotating” them. During annotation the sequences are analyzed and compared to existing sequence data in order to identify regions of the genetic code that we already know something about. Depending on the tools used and the scientific goals, annotation may identify the physical structure of the DNA, its functional properties (e.g. what proteins it produces), or even the organisms that are known to have this particular sequence. Annotation is both a computation- and data-intensive process. Successfully annotating sequences requires comprehensive and well-curated database of known sequences to which the new sequence can be

compared. And even the most powerful automated annotation systems require several hours to several days to annotate the data produced in a single run on a current DNA sequencing machine.

After annotation, researchers will analyze the annotated sequences using statistical analysis packages and visualization tools. While there are some “off-the-shelf” packages available, frequently these analyses are conducted using custom software and analysis scripts.

For all of these steps, but especially for annotation and certain forms of statistical analysis, the scientist must compare the sequences they are studying to other known sequence data. The need to assemble, collect, compare, and annotate large volumes of DNA sequence data precipitates numerous databases. How these databases are produced and used have implications for collaborative work and technology design.

The Ideal Database

An underlying theme across our participant interviews is the notion of an ideal Database.¹ Our participants talked about being able to share data, across what we in CSCW would describe as communities of practice, implying that existing databases are serving as boundary objects and consequently satisfying the informational requirements of all. Further investigation, however, shows that the notions of the Database are highly idealized and when delving into the details of practice, the successful use of these databases requires a great deal of translational and interpretive work.

Scientists using metagenomic approaches have a strong sense that sequence data is a public good.² The scientists we interviewed are keenly aware that to conduct research in this area it is necessary to have access to the data of others in order to compare genes at hand against previously found genes. In order to gauge environmental trends across time and space and to ascertain the unique qualities of particular genes requires access to amounts of data so vast that no single researcher or group of researchers could collect enough data. The “Database” is a particularly evocative concept here: a key feature of metagenomics research is that all prior sequence data serves as the baseline against which new sequences are compared.

A biologist working on the design of a database system described one scientific rationale for creating collections of sequence data:

¹ We use a capital “D” when we are referring to the conceptual ideal Database, and a small “d” when referring to a specific database system.

² It should be noted that most of our informants were working on government- or foundation-sponsored research in academic settings. However, even among those scientists who were involved with commercial research, we found that proprietary concerns might delay, but usually would not prevent the public release of sequence data.

In order to understand a new gene that we don't know what it does, we need to compare it with all the other genes that we have in the database that we know what they do. So we know what to do, for example, because they have been experimentally verified. We compare the sequences, computationally and we find the sequence in the database, and based on that, we can predict what the function of the gene may be. So then, what we also do is we integrate all that data in a single database because that is what is facilitating the comparative analysis; you must integrate all the data.

Comparing a new sequence against the Database can reveal the identity and function of genes and organisms. Similar comparisons to and analyses of “all the data” are used to understand the evolutionary history of organisms or the diversity of microbial populations. These comparisons are even useful for determining if a gene has been previously identified for patent applications.

One of the themes that emerged from our interviews and observations is that metagenomics researchers, bioinformaticists, and developers of cyberinfrastructure often have a strong sense of an ideal Database that they would like to have available or are actively trying to create. While the details of the ideal Database vary from person to person, generally it holds all available sequence data and associated metadata, and often derivative analyses. The Data would be well classified and annotated, and the Database would not contain errors or redundancies.

The ideal Database is also explicitly collaborative. The Data would be collected from and useful to scientists from a wide variety of communities of practice. Part of the rationale for spending the large sums of money required to develop such comprehensive databases is because the Data could reach across so many domains. The same sequence Data are potentially useful for medical research, environmental remediation, energy production, national security, drug development, chemical production, and many other pursuits. The Database is intended to be a boundary object, providing a standardized repository supporting cooperation across multiple communities of practice.

Our informants tend to think of the ideal Database as separate from the specific database systems they use in their work. One metagenomics researcher spoke of the Database this way:

We also rely on data that is in the database.... We generate a lot of primary data ourselves, but if we want to make comparisons, we have to compare to what's in the database. So we will use EMBL and GENBANK and the data that's in those databases as well.

This scientist refers to all of the sequence data produced outside of his laboratory as “*in the database.*” He then goes on to list specific sequence database systems he uses. This was a common trope across many of our interviews with scientists. When we asked them to describe the process of analyzing sequence data, they would often say that they compared the sequence data they generated in their laboratories to “all the other sequence data” or “every other known sequence.” On the other hand, when we observed their work or asked the scientists to tell us

about the specific databases they used, we found that they often used multiple databases, none of which actually contained all the Data.

The lack of integration of databases and datasets creates usability problems for scientists. As no single database contains all of the Data, scientists will often create their own local aggregated datasets to work with:

It's very important to have all metagenomes gathered together in one platform so that when people look for metagenomes they don't have to go here and here and, you know, it makes it all convenient... Yes, if you were doing a complex analysis and gathering data from many metagenomes and you would have to register on this server and this server and also this server to get the metagenomes. And then this server would provide you with some information, this other one slightly different information and the third one another kind of different information. It just makes it really hard if the data is all scattered around.

There are a number of projects that are working to make the ideal sequence Database a reality. Our respondents reported using many other database systems including GENBANK, the EUROPEAN MOLECULAR BIOLOGY LABORATORY NUCLEOTIDE SEQUENCE DATABASE (EMBL-BANK), the COMMUNITY CYBERINFRASTRUCTURE FOR ADVANCED MARINE MICROBIAL ECOLOGY RESEARCH AND ANALYSIS (CAMERA), THE SEED, INTEGRATED MICROBIAL GENOMES (IMG/M) and others. Probably the most well-known and longest-lived example is GENBANK, which was created in 1982, and is now housed at the National Center for Biotechnology Information, a division of the National Library of Medicine at the National Institutes of Health. GENBANK was founded to be "an annotated collection of all publicly available DNA sequences" (National Center for Biotechnology Information). A developer of another database system told us about the amount of effort spent to try to create a comprehensive database:

We're very aggressive in going out and getting basically all the data that relate to the public domain and integrate them. This is one of the most intense parts of maintaining and updating the system, constantly updating and adding everything that is released to the public domain.

This is made more difficult because sequence data are being generated in many locations, and even GENBANK does not contain all of the publicly available sequences:

There are several other sequencing centers that do not directly submit their data into GENBANK. They are keeping the data and releasing them through their websites, but they are not necessarily depositing them directly into GENBANK.

But even with the incompleteness of individual database instantiations, participants still expressed confidence in the ideal Database. One developer of a competing database systems told us:

It doesn't serve the community well if [our database] stands out there distinguished, beating its chest, saying we have more data than [GENBANK] or we have different data than [GENBANK]. I would argue philosophically that's a losing strategy and [our database] should not distinguish itself on what data it contains.

Many of our informants felt that what "the science" and "the community" required was for all of the specific sequence database systems to be operating on the same set of Data. Projects are underway to facilitate the creation of this universal

Dataset across database systems. For example, the GENOMIC ROSETTA STONE “is creating a mapping of identifiers describing complete genomes across a wide range of relevant databases so that information about genomes and the organism from which they derive can be more easily integrated” (Genomic Standards Consortium, 2008). Their vision is to create a distributed but easily accessible version of the universal Dataset by connecting many database systems into a federated database.

However committed scientists and database developers are to realizing a concrete version of this abstract ideal Database, we find that the vision for the Database is contested. Both the ideal Database and particular database systems are implicated in ongoing controversies about appropriate research questions, the role of the researcher, science funding, and scientific validity. At the same time that the Database supports collaboration, it is also playing a role in the active negotiation of practices and understandings. Rather than passing easily between communities of practice, using the databases requires significant translational work. Every scientist we spoke with reported using multiple databases, often having to manually reformat, edit, and combine the outputs of different databases. The databases often do not contain the contextual information necessary to make sense of the sequence data. Frequently this results in frustration for both users and developers.

The big issues with metagenomics is that the big archives are dysfunctional. They’re not only dysfunctional for metagenomics, they’re also dysfunctional for genomics these days.

The Database is intended to be a boundary object, but we believe that it is more productive to understand both the ideal Database and the individual instantiations as boundary negotiating artifacts.

The Database as Boundary Negotiating Artifact

Participants describe the individual sequence database systems as if they were shadows, poor representations of a widely-agreed-upon ideal. We find, however, that by looking across the landscape of databases, a different picture emerges. Instead, each decision about the implementation of a particular database system plants a stake for a community boundary. The databases are not so much imperfect copies of an ideal as they are arguments about what the ideal Database should be.

In this section, we will draw on our observations and interviews to discuss two areas of negotiation around the Database. First, we will discuss the close relationship between the Data and local scientific practice. Then we will discuss the problem of metadata and information completeness in sequence databases. In both instances, we find that rather than being stable boundary objects that move across community boundaries, the databases are malleable artifacts that serve as sites for negotiation of community boundaries.

The Database in Practice

The Database both contributes to and results from scientific practice. To understand this claim, it is important to look a little deeper at the technical implementation of what are commonly called databases. Our respondents speak of the databases as collections of data, but that only tells part of the story. It is more accurate to think of them as database-driven systems that include some combination of raw sequence data, contextual metadata (data about the environment from which the sample was collected), procedural metadata (data about how the samples were processed), assembled sequences, annotations, pre-computed analyses, and various tools for data comparison, annotation, visualization and analysis.

The particular arrangement of data and tools that make up each of these systems is driven by particular scientific needs. One scientist involved in database development told us:

So I kind of think it goes back to having that question, right.... What's your underlying emphasis for having the database? So, our underlying emphasis is that we have some questions that we're trying to answer both in complete genomes and for metagenomes.... Some of the things that we're trying to do is to take really specific problems that we're trying to address and use [our database system] to address some of those problems.

One of the ways that databases are tuned for particular research questions is through their accession policy. A typical strategy is for databases to focus on a particular type of organism or environment. For example, there are databases that focus on marine microorganisms, soil microorganisms, organisms found in the human gut, etc. Another database is attempting to collect data only about pathogenic organisms.

Another strategy is to focus on a particular type of data, regardless of the source. For example, some databases are collecting only “16S ribosomal RNA” sequences. These sequences are subunits of RNA that are useful for studying the evolutionary relatedness of species. But these sequences (which also appear in more general archives like GENBANK) are applicable only to specific kinds of research questions. One scientist who studies microbes that cause various diseases explained why she did not use 16S databases:

Everybody uses 16S and 18S sequences to categorize the phylogenetic community present. But similar organisms may have the same 16S, but have completely different physiology. So some, like *vibrio* for example, they're a great example of this. Many *vibrios* have the same 16S but can acquire a few genes, either by horizontal gene transfer, by phage transfer and they become highly virulent. *Vibrio cholerae* is a great example of that. You can have *vibrio cholerae* that's not toxic at all. It acquires one gene from its phage, the CTX gene. Horribly virulent organism, but if you look at the 16S, you'll never know.

In other words, the method used by many metagenomics researchers to categorize an organism, is useless for certain types of questions such as those about whether or not an organism is toxic or infectious. 16S sequence databases are useful for understanding how species relate to each other, but they are not sufficient for

understanding how variations in other parts of the genome can lead to functional differences in microorganisms. By choosing to only collect 16S sequences, the database developers have privileged certain scientific questions over others.

Database systems are also customized with particular query and analysis tools. One advisor to a database system told us about the problem of inheriting data and tools from a different research community:

So they had a lot of approaches to data analysis. Now, what they were looking for was slightly different than what the [microbial ecologists] were looking for, but it was something that many people were interested in. They were basically on a hunt for genes.... Eventually, all of the data [they] had in hand plus all of their analysis, eventually became the first datasets in [our] database. And many of the database structures in the database tools were developed by [them]. And for the purposes that they had at the time they did that, I think the database was actually adequate and not too bad. The problem was that it didn't serve [our] community quite as effectively as one would like. And so we made a series of recommendations over time about restructuring the system to be more accommodating to the kinds of questions that the ecologists were asking rather than the kinds of questions that molecular biologists and gene finders were asking.

When the microbial ecology project adopted the database system from the traditional genomic “gene finders,” they expected the database to be a boundary object. They knew they would have to customize it to some extent, but thought it would be able to “travel across borders and maintain some sort of constant identity” (Bowker, et al., 1999, p. 16). In the end, however, the system was so tailored to a specific set of research questions that the collection of data, the set of tools, and even the social organization of the project had to be significantly changed. New analysis tools were developed and old tools were discarded. Not only was the database ported to a different technology, the data itself was significantly restructured to fit the new tools and approaches. While the database development projects had begun by working together, in the end they were unable to collaborate. The system that was supposed to tie these groups together could not be shielded from the controversies that formed the boundaries between the communities of practice.

Metadata and Informational Needs

One of the features distinguishing metagenomic approaches from traditional genomics is a reliance on contextual data, or *metadata*. Unlike traditional genomics that focuses on the genetic information in a single organism, metagenomics considers the relationships of populations of microorganisms to their environments. In order to understand, for example, the effect of changing ocean temperatures on microbial populations, it is necessary not only to have sequence data but also to have associated data about where and when the sample was taken. Ideally, every sequence in the database would be linked to data about the environment from which the sample was taken, the people involved in the samples collection and processing, and the procedures used to isolate and

sequence the DNA. But collecting, storing, and disseminating this metadata adds another layer of complexity to the technical exercise of database development. Metadata and metadata standards become contested artifacts and sites of negotiation within the metagenomics and wider genomics communities.

One of the defining characteristics of boundary objects is that they are able to satisfy the information needs of different communities of practice. However, changes in the information needs of the community and the inclusion of new communities can challenge the ability of an artifact to be a boundary object. Metagenomics brings new questions, and existing sequence databases are inadequate for the metagenomics community's needs.

Until recently, most existing sequence databases had little, if any, metadata support. Even if scientists wanted to share metadata through the database, often their only recourse was to add a comment in a free-text field. More commonly, a scientist wanting to know more about a sequence in the database would have to track down associated publications and hope that the authors had included the relevant details. A program officer from a funding agency described the message coming from metagenomics researchers:

The community of principal investigators basically said, "Look, there's all these [metagenomic] data coming down.... The existing databases are simply not capable of providing us with the ability to do what we need to do with these data. You've got to do something about this. Because otherwise all of these data will be lost to us or to the scientific community because the ability to query on these data will just be gone. It won't happen if you don't do something."

Not only were there no metadata-capable databases, but there were no standards for what contextual data should be collected or how to represent it for storage. Scientists will typically only collect the data that is relevant for the study at hand. One scientist expressed frustration about the difficulty of sharing metadata:

You don't measure salinity when you work in the ocean. Right? You just assume the salinity is about the same.... Unless you've got a CTD [conductivity, temperature, and depth sensor] or something.... It really depends on what your question is. What I think is important as metadata, in fact what I know is important as metadata, nobody will ever measure.... We're doing microbial ecology. Essentially nobody measures what the microbes are eating.... That's because it's a hard thing to measure. But they'll all have nutrient analysis, though. That's because nutrient analyses are easy to measure.

Having metadata standards is important for both scientists and database developers. For scientists, a standard can function as a guide for what data to collect and how to represent it. For the database developers, the standard outlines what data should be in the database. Metadata standards are in active development, and some have even been published (Field, et al., 2008), but these standards are still being negotiated and none have been widely adopted.

But the adoption of these standards reveals the way that the Database not only crosses boundaries but is also implicated in pushing and establishing boundaries. Environmental metadata is extremely important for microbial ecologists, is less important for some other metagenomic questions, and is significantly less

important to many traditional gene- and whole-genome-focused users of the Database. The upshot is that it is important to the microbial ecologist that the geneticist attaches environmental metadata to sequence data, but it is not important to the geneticist. Similarly, the metadata needed by a marine microbiologist is significantly different from that needed by someone studying the microbial population of the human gut. This is a classic case of a disparity between those who must do extra work and those who benefit from the work (Grudin, 1989).

In the face of this difference in the value of metadata between communities of practice, the database becomes an important site for negotiations of the division of labor. In a discussion of metadata standards, developers of sequence databases were asked to require contributors to submit standards-compliant metadata with their sequence data. Databases that could not (or would not) make metadata a requirement were asked to alter the interface to make metadata submission easier, to make it easier to limit searches to sequences with metadata, and to create certification programs to give special status to sequences with compliant metadata. To use the language of Latour (Latour, 1987), the database becomes a mechanism for enrolling and controlling others in the creation of a particular kind of science.

Supporting Collaboration

While biologists, computational biologists, bioinformaticists, and computer scientists take the need to work together as given, collaborative endeavors differ according to content and scope. What matters are the particular scientific questions, not disciplinary allegiances or training. Each set of scientific concerns requires different types of metadata and different types of output. The database is a common denominator but is not sufficient for accomplishing work. Collaboration in the metagenomics area can be crudely classed according to whether they prioritize biological or environmental questions, but upon further investigation those classes quickly breakdown into subcategories with some overlapping and some unique requirements.

We found that the fit between the database system and the scientists' research questions was a more important decision factor in choosing a database to use than the completeness of the database. One researcher told us about certain databases being better repositories than others, but then when asked why he chose to use particular databases, he said:

Researcher: Because of what it does. Because of what I can get out of it.

Interviewer: Is it about the tools or the data that they have?

Researcher: It's about the results. The different websites - you can get the data from any of those websites. It's about the tools. It's about the results that they can produce for you.

Another researcher also emphasized the importance of the visualization and analysis tools:

Well, originally I started using [that] database because it's a great way to look at functional analysis.... What you're looking at is at the functional profiles of each one of these samples. And like I said before, that's really important for understanding the function of a community.... And so we can track - and perhaps if you think of an ecosystem, or in this case, a metabolic system, looking at shifts in the metabolism of the whole system might be more environmentally relevant than just looking at the change in a particular strain of bacteria. And so you can see these massive changes; all sorts of great ways to parse the data into something that's biologically relevant.

For this scientist, the most important criteria for choosing this database was the ability to analyze and visualize the data in a way that made it “biologically relevant” to the questions she wanted to answer. For these scientists, the best database was not the one that came closest to the ideal comprehensive “all known sequences” Database. Instead, it was the one that best fit the research, in other words, the one where the entire database system—data, structures, tools, and outputs—came together to best support the scientist’s practice and produce the most meaningful answers to scientific questions.

Designing databases that work for scientists is an immense challenge. There is a great diversity of need regarding data, metadata, and software tools that stems from a diversity of scientists and scientific interests. Advances have been made with technologies like ontologies, which can provide semantic mappings across domains (Schuurman & Leszczynski, 2008). But the challenges described here go beyond semantics and invoke questions of the value and organization of scientific practice. A metagenomicist who has collaborated in the development of a database system told us about three stakeholder communities that are trying to use a particular database. He describes the groups as moving targets:

There are at least three moving targets in this project. And that is that there are the ecologist metagenomics people, there are evolution people that are more interested in the evolution of the sequences, you know, what they're telling you about evolution; which is actually quite different how you analyze the data in this case. And then there are just the people that are thinking, like, just genomes and glorified genomes, right. And that's also a very different way of looking at the world. I think that that's a big failing that we didn't recognize that in the beginning as much as we should have....

We need different outputs. That's kind of the problem. So you do almost the same thing in the beginning, but if you're interested in the genome, you want a genome browser, right, you want to scroll on a genome and look at a gene, where it's at and everything. If you're someone like me, you want something that can be funneled into a statistical package. And if you're an evolutionist, you want the same information, but you want to be able to do an alignment with them. I mean it's the same exact analysis, but a different 'what do you do with it at the end' sort of thing.... It's the tools that really count.

What we see is that there are a number of communities using these database systems, and each brings its own set of research questions and viewpoints. At a base level they are all using the “same” sequence data. But in practice, the Data do not exist independently of the database system. Even accessing “raw” data

requires understanding the particular data collection in the given system (along with its potential errors, omissions, and redundancies), navigating a particular set of data standards and formats, and dealing with the particular query and output technologies. When a system lets you select data based on the presence of a particular gene but not based on the geographic location from which they were sampled, that system reinforces a particular set of research questions and strengthens the boundaries between communities of practice.

While there is a diversity of scientists and scientific interests, the larger challenge is that the research questions are continually evolving.

The instruments will continue to improve. But they're never going to be perfect because we're continuing to push the boundaries. So the kinds of scientific questions we can answer will keep extending. So we'll have demands for new instrumentation. We'll have demands for new software tools. And I can make up 20 questions that are important today. Any microecologist can make up better questions, probably, than I can about really - some the same, some different and better than I can about the challenges; the new questions that metagenomics will allow us to ask and answer. But I think we also feel that we don't know the range of questions fully. And so the same is true for software tools.

As scientists are successful at generating new discoveries, as technologists are successful at developing new technologies, and as these innovations synergistically drive each other forward, the research questions will change and the range of answerable research questions will also change. Many of our participants discuss being drawn to "metagenomics tools" and some of our participants refer to metagenomics as a new discipline. It is not a far stretch then to assume that as research questions, data, tools, and practices shift and change so too will the communities around them.

Those wishing to support scientific collaboration should take care to map out scientific stakeholders according to scientific questions, and not according to domain or institutional allegiances. This mapping out of concerns must be done iteratively to keep pace with scientific developments. Furthermore, CSCW researchers in this area should be aware that scientists will talk about databases as if they function as boundary objects, but that when pressed for more detail, scientists reveal that their databases require a great deal of work in order to meet the needs of different communities of practice. This latter phenomenon is not necessarily a failure of requirements specification. The requirements are often sufficient at the time of collection but are rendered inadequate by scientific advances. Furthermore the multiplicity of databases is a reflection of the multiplicity of interests and competing knowledge claims that are indicative of a vigorous scientific community. Some degree of integration may be desirable, even inevitable, but a high degree of integration among research databases is a mirage—a utopian ideal. A lofty goal for computer supported cooperative science would be to find ways to support simultaneously cooperative data sharing and scientifically competitive (i.e. divergent or unique) data acquisition use, analysis, and theory building.

Conclusion

New types of science also require new standards, processes, and collaborative social structures, such as distributed virtual organizations comprised of domain scientists, information scientists, and engineers. Metagenomic science is among many endeavors that require work to be coordinated through and around multiple databases. More work is needed to understand how collaborative work is structured by multiple databases. The tendency to dismiss situations where organizations depend on imperfectly interoperable databases as merely inefficient legacy systems is likely glossing over insights about just how multiple databases support not only different types of work but also different perspectives and priorities. There is a more interesting story to tell about how they actually serve to support and constrain work. We also find that there is an important connection between multiple databases and coordinative artifacts.

As mentioned in earlier research on boundary negotiating artifacts in a small, nascent design group using primarily paper documents, artifacts can be used to cross boundaries between communities of practice. But they can also be used to affect the division of labor, or in other words, to push and establish the boundaries between communities of practice (Lee, 2007). In this paper we have looked at metagenomic science, which is a very different sort of endeavor, and yet it too requires complex coordination around another type of artifact: the database. Rather than looking at databases as static, we choose to look at databases as existing across scientific communities where the scientists involved have different practices and priorities. Due to rapid scientific and technical innovation the tools, practices, and scientific questions change over the course of merely a few years. The sequences within the databases are relatively static, being a sort of minimum common denominator, but what make the databases useful and relevant are the array of constantly-changing software tools and highly negotiated metadata.

In dynamic environments, the number of true boundary objects that satisfy the information requirements of multiple communities of practice may be relatively few compared to the number of prototypes, boundary objects in-the-making, or boundary negotiating artifacts. If we can consider the database to be another type of artifact that coordinates multiple perspectives, we begin to see how multiple databases may sometimes be necessary and useful. The challenge for computer supported cooperative science then becomes how to meaningfully support large-scale collaborations that are reliant on multiple databases that support a multiplicity of knowledge building priorities and practices.

Acknowledgments

Special thanks to the participants who were so generous with their time, Eric Baumer, who assisted in data collection, and the anonymous reviewers for their comments. This work was supported by National Science Foundation awards IIS-0712994 and OCI-083860.

References

- Atkins, D. E., Droegemeier, K. K., Feldman, S. I., Garcia-Molina, H., Klein, M. L., Messina, P., et al. (2003): *Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel On Cyberinfrastructure*. Washington, D.C.: National Science Foundation.
- Birnholtz, J., & Bietz, M. J. (2003): 'Data at work: Supporting sharing in science and engineering' *Proceedings of the 2003 International ACM SIGGROUP Conference on Supporting Group Work*, New York, NY: ACM Press, pp. 339-348.
- Bowker, G. C., & Star, S. L. (1999): *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press.
- Chen, K., & Pachter, L. (2005): 'Bioinformatics for whole-genome shotgun sequencing of microbial communities', *PLoS Computational Biology*, vol. 1, no. 2, Jul, pp. 106-112.
- Emerson, R. M., Fretz, R. I., & Shaw, L. L. (1995): *Writing Ethnographic Fieldnotes*. Chicago, IL: University of Chicago Press.
- Field, D., Garrity, G., Gray, T., Morrison, N., Selengut, J., Sterk, P., et al. (2008): 'The minimum information about a genome sequence (MIGS) specification', *Nature Biotechnology*, vol. 26, no. 5, May 2008, pp. 541-547.
- Genomic Standards Consortium (2008): 'Genomic Rosetta Stone', Retrieved March 5, 2009, from http://gensc.org/gc_wiki/index.php/Genomic_Rosetta_Stone
- Gilbert, J. A., Thomas, S., Cooley, N. A., Kulakova, A., Field, D., Booth, T., et al. (2009): 'Potential for phosphonacetate utilization by marine bacteria in temperate coastal waters', *Environmental Microbiology*, vol. 11, no. 1, Jan, pp. 111-125.
- Glaser, B. G., & Strauss, A. L. (1967): *The Discovery of Grounded Theory: Strategies for Qualitative Research*. New York: Aldine de Gruyter.
- Grudin, J. (1989): 'Why groupware applications fail: problems in design and evaluation', *Office: Technology and People*, vol. 4, no. 3, pp. 245-264.
- Handelsman, J., Rondon, M. R., Brady, S. F., Clardy, J., & Goodman, R. M. (1998): 'Molecular biological access to the chemistry of unknown soil microbes: a new frontier for natural products', *Chemical Biology*, vol. 5, no. 10, Oct, pp. R245-249.
- Harper, R. (1998): *Inside the IMF: An Ethnography of Documents, Technology and Organizational Action*. San Diego: Academic Press.
- Harper, R., Procter, R., Randall, D., & Rouncefield (2001): "Safety in numbers": Calculation and document re-use in knowledge work' *Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work*, New York: ACM, pp. 242-251.
- Henderson, K. (1999): *On Line and On Paper: Visual Representations, Visual Culture, and Computer Graphics in Design Engineering*. Cambridge, MA: MIT Press.
- Hilgartner, S. (1995): 'Biomolecular databases: New communication regimes for biology?', *Science Communication*, vol. 17, no. 2, pp. 240-263.
- Hine, C. (2006): 'Databases as scientific instruments and their role in the ordering of scientific work', *Social Studies of Science*, vol. 36, no. 2, April 1, 2006, pp. 269-298.

- Hugenholtz, P., Goebel, B. M., & Pace, N. R. (1998): 'Impact of culture-independent studies on the emerging phylogenetic view of bacterial diversity', *Journal of Bacteriology*, vol. 180, no. 18, pp. 4765-4774.
- Latour, B. (1987): *Science in Action*. Cambridge, MA: Harvard University Press.
- Lee, C. P. (2007): 'Boundary negotiating artifacts: Unbinding the routine of boundary objects and embracing chaos in collaborative work', *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, vol. 16, no. 3, pp. 307-339.
- Lutters, W. G., & Ackerman, M. S. (2002): 'Achieving safety: A field study of boundary objects in aircraft technical support' *Proceedings of the 2002 ACM Conference on Computer Supported Cooperative Work*, New York: ACM, pp. 266-275.
- Manovich, L. (2001): *The Language of New Media*. Cambridge: MIT Press.
- National Center for Biotechnology Information (April 2, 2008): 'GenBank Overview', Retrieved February 23, 2009, from <http://www.ncbi.nlm.nih.gov/Genbank/index.html>
- O'Day, V., Adler, A., Kuchinsky, A., & Bouch, A. (2001): 'When worlds collide: Molecular biology as interdisciplinary collaboration', in W. Prinz, M. Jarke, Y. Rogers, K. Schmidt & V. Wulf (eds.), *Proceedings of the Seventh European Conference on Computer-Supported Cooperative Work*, Dordrecht, Netherlands: Kluwer, pp. 399-418.
- Pawlowski, S. D., Robey, D., & Raven, A. (2000): 'Supporting shared information systems: Boundary objects, communities, and brokering' *Proceedings of the 21st International Conference on Information Systems*, Atlanta, GA: Association for Information Systems, pp. 329-338.
- Schmidt, K., & Simone, C. (1996): 'Coordination mechanisms: Towards a conceptual foundation of CSCW systems design', *Computer Supported Cooperative Work (CSCW)*, vol. 5, no. 2, pp. 155-200.
- Schmidt, K., & Wagner, I. (2002): 'Coordinative artifacts in architectural practice', in M. Blay-Fornarino, A. M. Pinna-Dery, K. Schmidt & I. Wagner (eds.), *Cooperative Systems Design: A Challenge of the Mobility Age*, Amsterdam, The Netherlands: IOS Press, pp. 257-274.
- Schmidt, K., & Wagner, I. (2005): 'Ordering systems: Coordinative practices and artifacts in architectural design and planning', *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, vol. 13, no. 5-6, pp. 349-408.
- Schuurman, N., & Leszczynski, A. (2008): 'Ontologies for bioinformatics', *Bioinformatics and Biology Insights*, vol. 2008, no. 2, pp. 187-200.
- Star, S. L. (1987-1989): 'The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving', in L. Gasser & M. N. Huhns (eds.), *Distributed Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann, Vol. II, pp. 37-54.
- Star, S. L., & Griesemer, J. R. (1989): 'Institutional ecology, 'translations' and boundary objects: amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39', *Social Studies of Science*, vol. 19, no. 3, pp. 387-420.
- Subrahmanian, E., Monarch, I., Konda, S., Granger, H., Milliken, R., Westerberg, A., et al. (2003): 'Boundary objects and prototypes at the interfaces of engineering design', *Computer Supported Cooperative Work: The Journal of Collaborative Computing*, vol. 12, no. 2, 2003, pp. 185-203.
- Turnbaugh, P. J., Backhed, F., Fulton, L., & Gordon, J. I. (2008): 'Diet-induced obesity is linked to marked but reversible alterations in the mouse distal gut microbiome', *Cell Host Microbe*, vol. 3, no. 4, Apr 17, pp. 213-223.
- Wenger, E. (1998): *Communities of Practice: Learning, Meaning, and Identity*. New York: Cambridge University Press.

Towards *LivingAgendas* – Shaping the next generation of business meetings

Till Schümmer¹, Hilda Tellioglu², Jörg M. Haake³
FernUniversität in Hagen^{1,3}, Vienna University of Technology²
*till.schuemmer@fernuni-hagen.de*¹
*hilda.tellioglu@tuwien.ac.at*²
*joerg.haake@fernuni-hagen.de*³

Abstract. Business meetings are omnipresent in all kinds of organizations. This paper presents an analysis of meetings at one specific medium-sized enterprise. By means of ethnographic studies, we observed collaboration and coordination problems in meetings. We address these problems with socio-technical meeting patterns, as documentations of good practices that help to understand and change the social interaction, the infrastructure, or both. These pattern-driven interventions helped us to gain insights into the socio-technical aspects of meetings. Finally, we created a first prototype of an integrated meeting support system.

Introduction

Virtual organizations are becoming more important these days. A virtual organization is “a temporary network of independent companies linked by the free flow of information” (Byrne et al., 1993). Meetings between collaborating members of different companies as well as meetings within individual partner organizations become an essential part for performing the business of the virtual organization. While we reported on cross-company meetings in previous work (Schümmer & Haake, 2009), this paper takes a closer look at meeting interactions at one specific company. We report on our observations of real-live business meetings at a medium-sized automotive supplier who provides services for several customer companies.

We first observed meetings using ethnographic methods and analyzed problems that arose from ignoring important properties of well-designed and well-organized meetings. We created a meeting pattern language containing descriptions of good practices. Patterns formed the basis for interventions, both at the level of group processes and at the level of supporting technology. Our interventions were triggered in a workshop in which the practitioners were introduced to the patterns. Our analysis focused on how social interaction would change when the observed groups started to apply the patterns on a social level while being supported by wiki technology. The feedback of the participants suggested that knowledge of the patterns forms a good starting point for improving meeting practices and for making informed use of groupware technology.

However, our users requested that standard technology could be improved so that there are fewer breakdowns caused by transitions between meetings, tools, and media. Based on this feedback, we created *LivingAgendas*, a first prototype of an integrated meeting support system.

Business meetings in virtual organizations

The question of how meeting quality can be improved has been subject of numerous, often management-oriented, textbooks and articles (e.g., Doyle & Strauss, 1976; Jay, 1993; Streibel, 2003; Kelsey & Plumb, 2004; Parker & Hoffman, 2006; Matson, 1996). Not surprisingly, all these books address comparable problems and come to comparable suggestions on how to improve face-to-face meetings. Although the advice often seems simple and straightforward, people still face difficulties following it in their daily meeting activities. In the research of Group Support Systems (GSS), meetings have been studied (e.g., de Vreede et al., 2002, Streitz et al., 2001), for instance, with the focus on the opportunity of equal participation during meetings, on the effectiveness and efficiency of meetings, on the quality of meeting results, and on the design of meeting room layouts, to name only some issues. Considering these guidelines and scientific studies, we derived ten fundamental properties of an efficient meeting:

Motivation and Reliability. Participants do not take the meeting seriously. Meetings are not considered as work. Participants arrive late or leave early and start to doodle.

Clear Goals. The meeting lacks an agenda and there is no concrete vision for an output or result of the meeting.

Information. Important information is not available in a meeting. Participants did not prepare in advance and thus lack important background information. As a result, no decisions can be made.

Context. The environment in which the meeting is held is inappropriate for the meeting's topic. The meeting is scheduled at the wrong point of time.

Focus and Efficiency. The participants gradually shift topics until they discuss issues that are totally unrelated to the agenda. They perform other tasks in parallel, which takes most of their attention.

Trust and Openness. Participants tell lies in the meeting. There are long discussions but no honest contributions. Conflict is avoided instead of being resolved.

Respect. Participants personally attack other participants. They start to look for mistakes made by others to make them lose their face.

Communication. Participants do not listen to one another. Too many participants speak at the same time.

Participation. Participants do not participate. Only few people engage in a discussion.

Results. There is no action after the meeting. Participants do not manage to implement decisions made. At the next meeting, resolved issues are discussed again.

Technology support becomes an important factor in virtual organizations since not all participants of a meeting are working at the same location. Even if meetings take place at the same location, the preparation is usually a distributed activity. One of the earliest works that investigated the role of technology support during the whole meeting life cycle is the *GroupSystems* study by Nunamaker and others (1991). They argued that an appropriate combination of group interaction tools could support processes and tasks as well as help to create a better structure for work processes. While these tools provided first clues towards integrated meeting support solutions, future developments mainly focused on specific phases of the meeting. Based on agent technologies, there are some systems for meeting scheduling for arranging meetings, some with more attention to user preferences (Jennings et al., 2003; Herlea et al., 2001). Bicharra Garcia et al. (2004) investigated mechanisms for increasing the quality of a meeting agenda, including the process of prioritizing agenda items suggested by prospective meeting participants. Other systems focus on in-meeting support. The system proposed by Vivacqua et al. (2008) makes use of scripts for creative facilitation, so-called *ThinkLets* (Kolfshoten et al., 2004). *ThinkLets* are tool-centered. They contain codified scripts as facilitation routines to execute actions and instructions. While *ThinkLets* already suggest a specific flow of interaction, they still require the practitioner to interpret the effect towards meeting properties. They also provide only limited support for embedding the interaction flow in an existing environment. The *Meeting Central* application (Yankelovich, 2004) focuses on sharing information and establishing speaker awareness within a meeting. Some systems focus on participation and involvement of all participants in a meeting: by means of interactive whiteboard applications or large computer screens, implicit interaction during co-located, ad-hoc meetings can be supported (Ju et al.,

2008; Smart Board, 2009). These systems enable, on the one hand, the visualization of shared information to all participants. On the other hand, they offer functions for direct modification of the presented content. *Roomware* augments the concept of interactive whiteboards to provide a coherent interaction experience using interactive and networked furniture (Streitz et al., 2007; Tandler, 2008). But again, these approaches do not offer sufficient support for the whole meeting life cycle. Systems like *Microsoft SharePoint* (2009) allow the collection of relevant information at a single place. They, e.g., list meeting objectives, attendees, the agenda, and relevant documents. They offer the possibility to create action items during meetings. However, they do not consider the collaboration processes involved in meeting preparation and execution.

Ethnographic evidences for the absence of properties

During an ethnographic study in MAPPER (IST-016527) we have observed the lack of the meeting properties in real life settings. While the project's main goal was the improvement of design and manufacturing processes, the ethnographic material also provided insights into how employees interacted in meetings.

In this paper we focus on *Alpha*, a supplier in the automotive sector. It produces several components like seat climate and motion control, head restraints, control cables, or gearshifts. It cooperates with several other suppliers where it is responsible for the project management. One way of dealing with interdependencies is using meetings as an arena for exchange, coordination, and planning. Empirical material about work practices was collected during two field visits in November 2005 and March 2006. During our first visit we were able to observe how projects are managed. We followed co-located and distributed meetings, project meetings as well as design reviews, and ongoing work at various workplaces in design, testing, and purchase. During our second visit we focused on the practitioners' interactions with the external suppliers and on the company's ways of managing projects, especially with a high level of innovation.

While the full set of work processes was documented elsewhere (Jacucci et al, 2006), this paper concentrates on observations made in meetings at *Alpha*. In this section we will first describe the regular project meetings and problems of current practices connected to these meetings. Then, we will show how participants get involved in meetings. Our main goal of this analysis is to identify, study, and analyze missing or weak properties in these settings.

Regular project meetings at Alpha with several to-do lists

Managers at *Alpha* called for weekly project meetings with the complete project team, no matter whether there were relevant issues for each participant. Some participants reported that it shows a lack of *respect* if they are invited to these

meetings and forced to be there for one or two hours without making any contributions, others said that the *goal* of the meeting is not clear and they do not know why they have to be present. This makes an active *participation* in the meeting almost impossible, also because of a lack of required *information* for preparing oneself before entering the meeting.

Sometimes project managers call for additional meetings. The main goal of these additional meetings is to discuss open issues or gather missing information for certain issues, mainly to prepare decision-making processes. These additional meetings are not in the scope of this paper.

Regular meetings are mainly seen as an arena for exchange, articulation and clarification. They should help to reach a shared understanding of the project's status and as a result create strategies and commitments for future steps. However, we could observe several issues in these meetings, which could be improved in many senses. First of all, only project managers are allowed to call these meetings and moderate them. Even if some project members did have, e.g., problems to carry out their tasks between the meetings or if they had certain important questions needed to be discussed in the project team, they did not have the possibility to call for an emergency meeting. While project managers were highly *motivated* in holding the meetings, project members would have needed meetings at other points in time. While project members were forced to respect the project manager's need for a meeting, the project members asserted a lack of *respect* for their individual needs.

Project meetings help project managers to assess the progress of the project, to clarify uncertainties, to define responsibilities, to set deadlines, to negotiate objectives, and to define new tasks. The main tool for orchestrating the setting is a *to-do list* (Figure 1): "As a project manager you are not anyone's boss, you cannot give orders, to-dos are a way of giving indirect orders, setting responsibilities and deadlines". The to-do lists are used as *meeting agendas*. They are stored as spreadsheets or text documents. Each line contains an (open) issue with responsible persons, deadlines, and status. Sometimes issues are grouped in categories to structure the subjects to deal with. The main problem participants have with the project manager's to-do list is that they are not informed about the content of these lists before the meeting. Except of items they know because they took their own notes during the last meeting, they are uncertain whether there will be new items which they personally will be made responsible for and whether they need to prepare information to discuss specific items in the agenda.

The effects of this lack of *information* are crucial: team members have no clear understanding of the meetings' goals since they do not know details about what will be discussed, they cannot answer issues arising in the meeting in an *efficient* way since they had no opportunity to prepare required information, and they may hesitate to provide open answers (*trust and openness*) since they were not able to discuss the answers with colleagues before.

The restrictive access to the manager's to-do list has also effects on the work between the meetings. Since not even a read access is permitted for project members, they cannot make use of the meeting *results* captured by the project manager. The only shared access to the to-do-list takes place during the meeting where it is projected to the wall. In a way, this supports the *communication* between all participants. However, the missing meeting minutes after the meeting causes an inefficient way of organizing and managing ongoing project activities.

R & D. testing			
Week.no	Issue	Resp.	Deadline
38.1	D-FMEA to be performed, component level. Meeting scheduled w54S.	TOEK	w 539
40.3	Verify styling on L-shape carrier due to open slots for brace. Meeting with customer w 544.4.	ANKV	w 42
44.1	Track outstanding issues in TR.	TOEK	Follow up

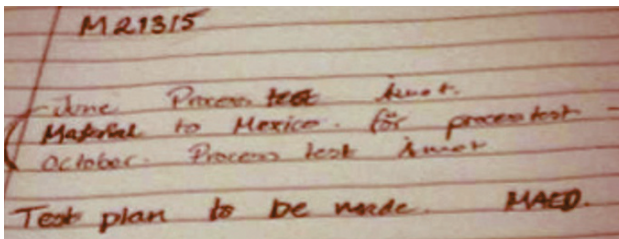


Figure 1. Computer-based to-do lists of project managers with general issues, responsible persons and deadlines (top) vs. handwritten to-do lists of meeting participants for individual use (bottom).

Having privileged access to the to-do list makes project managers special: They expect to have an overview of open issues and the status of work progress. However, this is only the case if they are informed enough. Work done in between the meetings is not transparent for them and is not reflected in the managers' to-do lists. This can be a problem, especially if they have to report to their managers. Sometimes the data is not up to date: project members have sometimes solved several open issues, but project managers are not informed, because there was no project meeting in the meantime. Such situations occur often and are not well seen in the company, especially if a customer acts as a partner in the project and is aware of inconsistencies in coordination and information flow in *Alpha*.

In addition to the project manager's to-do-list, participants note down their *individual to-do lists* during the meetings (Figure 1). This is a necessity for the team members. They need to structure and prioritize their ongoing work, to highlight their deadlines, questions they have to answer, interdependencies they have to consider, to plan and organize their communication with others in the project or with external partners, to remind the documents they have to create or update, etc. Their only possibility to gather this information is making their own notes during the meetings.

There are some inconsistencies caused by the fact that there is *no common document or information space* where each project member has access to modify issues he or she is responsible for. Each time project members start with a new individual list. There is no link to previous lists, which makes it more difficult to focus on *results* and use these as *information* for the next meeting. Sometimes they forget things, sometimes there are misunderstandings, and sometimes they carry out additional work without considering the work of colleagues because there is no shared awareness of current *results* and agreements.

We did not observe how they document their activities between meetings, whether they add their remarks to their notes from the meeting or they create other documents not attached to their meeting notes or whether they did not document their activities at all. What we could observe is that there was no additional written document or an electronic space to access during the meetings, in order to see this type of *information*. It was mainly a verbal reporting to the question of the project manager. If a project member could not join the meeting, which is a sign of lack of *participation*, there was no way to find out what he or she has done, what is still to do, where the problems were, etc. This caused problems regarding *reliability*.

The dynamics of the meetings: Interactions, artifacts, conventions

Meetings at *Alpha* have their dynamics: participants interact, several artifacts are used, and articulation work is supported by mockups or prototypes. There are different levels of *participation*: People come and go, talk or stay quietly, are active or passive. We want to show such a setting by the following illustration. It is a meeting in a customer project for seat ventilation, a project in its very late stage, with production scheduled to start early the following year. Participants arrive with their notebooks and calendars and create their *context*. They setup their individual meeting environment before the meeting starts. One had brought the new prototype of the head restraint (Figure 2), which he then uses to indicate design changes and demonstrate what had been tested. This relevant *information* is brought in during the meeting, but in this case, it is bound to a real artifact that is difficult to distribute before the meeting.



Figure 2. Prototype of head restraint used during the meeting.

The project manager opens his updated to-do list, which is projected onto the wall and is used for organizing the meeting and for noting decisions. Notes taken are made visible for participants, however only during the meeting (analogue to the previous case). While they discuss the issues in the sequence given by the list, the project manager writes directly into the document, adding or changing text and highlights particular entries, using color (red), bold type, and language such as “very urgent”. This creates a shared *focus* during the meeting, but the *results* are still not shared afterwards.

There is a sequence in the flow of talking: First, the project manager starts talking by reading the issue from his to-do list, mentioning the open questions or decisions made so far, asking the responsible person, whether the task has been closed or whether there are problems, etc. This activates the person addressed by him. He or she looks for documents or tries to read from his or her computer before answering the question. Sometimes he or she uses an artifact to explain why he or she changed the course of action since the last meeting. It is important to note that the participants cannot prepare answers. *Information* may be unavailable with the effect that the addressed person cannot provide any answers.

For instance, the project manager asks the participants: “Did AA [the customer] accept the soft tool quote? – Commercially there has been no feedback so far, the updated quote was sent 21st of October.” First, no one answers this question. No one feels addressed by this question – again a question of *participation* and pre-established awareness of *goals*. There is an unpleasant quietness in the room. Obviously, the participants who fear to lose their faces consider this situation as dangerous. The managers were not aware of these feelings and ignored them instead of *respecting* individual uncertainties. The project manager repeats his question. One participant looks first to his colleagues in the room and then tries to answer it. The project manager is not satisfied by the answer and marks this issue as “very urgent” and colors it red. Now, he puts the name of this participant as the responsible person for this item without asking for his commitment. Although a *result* was noted in the manager’s to-do-list, there was no open discussion on how to proceed with this issue (*trust and openness*). In this sense, management ignored a potential conflict in the team.

After a while the project manager reports on a steering committee meeting scheduled at the same day where he is forced to send a clear reminder to the customer. This is obviously a critical issue for him, and it could have been avoided if the *information* was gathered during the preparation phase of the observed meeting.

The example shows how issues are dealt with, step-by-step, quite quickly, and only sometimes a discussion comes up. Participants rarely take an active role. We also saw that physical artifacts, such as materials or prototypes have an important role in these meetings.

We are not so much interested in the details of participants' negotiations here than in the meeting dynamics, which can be characterized as document driven – as we see the project manager's to-do list is the central but not the only document used and referred to in this meeting. Each participant creates his or her own list of open issues which he or she is responsible for. However, there is no shared *information* space.

Each issue ends up in a *result* – a new task, the reformulation of a task, confirmed or modified deadlines, and so forth. But the decisions are badly *communicated* in the group and most important not accessible for the group members after the meeting.

Participants act as owners of particular issues. They enter and leave the discussion as these issues arise. Sometimes, they leave the room physically for a couple of minutes based on a personal judgment of the timing of upcoming issues in the to-do list (*participation*).

Many issues indicate uncertainties that have to be settled as fast as possible – a customer order for particular design changes to be obtained, a specific test to be scheduled, a missing document to be retrieved, a knowledgeable person to be contacted, and so forth (*information, efficiency*).

Although there is no overview of the project history present in the meeting, some of it is present in the form of issues and tasks that have been formulated and agreed upon but not yet resolved and for which eventually new deadlines have to be defined. It would on the one hand support decision-making in the group and avoid repeating same discussions in the course of projects, and on the other hand contribute to *clear goals*, if changes to single items and to the project as a whole could be made visible.

Meeting patterns: a tool for socio-technical interventions in meetings

As we saw from the ethnographic studies, there are both technical and social factors that make meetings suboptimal. For instance, the restricted access to the to-do list both has a social aspect, i.e. that the manager defines his role through a surplus of information, and a technical aspect, i.e. that the used technology does not allow concurrent access to the to-do-lists.

In order to structure our interventions for improving meetings at *Alpha*, we created a meeting pattern language. It is a collection of good practices for meetings that gives the project team advice on how to improve their meetings. The idea of using patterns for studying and improving socio-technical systems has been discussed by numerous authors before (e.g., Schümmer & Lukosch, 2007; Schümmer, 2005; Herrmann et al., 2003; Guy, 2005; Carroll & Farooq 2007). Basically, all these approaches have their origins in the work of the architect

Christopher Alexander (1977), who had the vision that patterns could empower lay people to act like construction experts and thereby allow them to change their buildings and cities.

The patterns of Alexander can already be considered as socio-technical patterns. They contain aspects that change the way people interact (the social component of the pattern) and aspects that change the environment in which they interact (the technical component of the pattern). A random example is the ADVENTURE PLAYGROUND pattern (Alexander et al., 1977). Here, the authors analyzed the social interaction between children on a playground and stated the problem that “any kind of playground which disturbs, or reduces, the role of imagination and makes the child more passive, more the recipient of someone else’s imagination, may look nice, may be clean, may be safe, may be healthy – but it just cannot satisfy the fundamental need which play is all about” (p.368f). The authors further discuss the need for adventurous and imaginative play and conclude by giving advice to the designers of playgrounds (especially the people living in the neighborhood who act as designers of their environment): “Set up a playground for the children in each neighborhood. Not a highly finished playground, with asphalt and swings, but a place with raw materials of all kinds ...” This technical part of the solution is concluded with a sentence that explains the intended social interaction in the space “... where children can create and re-create playgrounds on their own.”

This combination of social and technical aspects in the solution of a pattern is very valuable when addressing changes in meetings what we applied in our case to empower our users. For that reason, our patterns make the different aspects of the solution explicit by distinguishing between

- a solution part that addresses primarily the *social interaction*,
- a solution part that employs *standard technology* that is widely available in the organization such as electronic mail, instant messaging systems, or wiki systems, and
- a solution part that informs designers of *integrated groupware* applications. This solution part makes use of patterns for computer-mediated interaction (Schümmer & Lukosch, 2007).

In this sense, our patterns can be used to bridge the gap between ethnographic observations and socio-technical interventions. This is closely related to the approach taken by Guy (2005) who presented a study where patterns were used to integrate use, evaluation, and design of collaborative systems. In our work, this connection is further underlined by relating the patterns to the meeting properties identified before. The absence of a property can guide the designing user to an appropriate pattern addressing the property.

In order to illustrate the pattern format used in our pattern language, we will now present an example pattern that addresses one of the most critical problems in the observed meetings at *Alpha*. It is a pattern that aims on making the agenda

a shared document. Note that pattern names will be shown in SMALL CAPS. Pattern followed by an asterisk can be found in (Schümmer & Lukosch, 2007).

An example pattern: IT'S MY AGENDA – IT'S MY MEETING

Context: You are calling for a meeting and create an agenda for it. There are stakeholders who have different backgrounds and interests.

Problem: The owner of the meeting normally creates an agenda. All other invited participants have only limited possibilities to participate in the agenda creation. This can lead to incomplete or wrong agendas.

Properties: Motivation, clear goals, and participation.

Symptoms: Apply the pattern to avoid the following to happen...

- Agreeing on the agenda takes a lot of time at the start of the meeting.
- Many new topics pop up during the meeting, which have not been foreseen.
- Not all interests are reflected in the agenda. People do not participate in the meeting, as it does not address their needs.
- Important topics do not make it to the agenda early enough for allowing good preparation of the topics.

Social Solution: Define a shared place where all invited meeting participants can collaboratively prepare the meeting agenda up to a specified deadline. Use this shared place also to collect input documents and presentations so that everyone can BE PREPARED. Mark those items in the proposed agenda that do not require face-to-face discussions in the meeting in order to reduce meeting time (see PRESENT WITHOUT PRESENTING and NO DISCUSSION). Creating the meeting agenda collaboratively helps creating a collaborative responsibility for the success of the meeting.

Standard Technology Solution: Before announcing the meeting, the meeting owner creates a wiki page with an agenda skeleton. For regular meetings, this skeleton contains entries for recurring agenda items. In the announcement message (sent by electronic mail), the organizer invites the participants to extend and/or modify the agenda. After changing the agenda, the participant informs all other participants by electronic mail that the agenda was changed. An easy way to do this is to reply to the initial invitation including all initial recipients as receivers of the notification mail. Shortly before the meeting, the meeting organizer sends a request for a VOTE* on the agenda.

Integrated Groupware Solution: The agenda and all agenda items are stored as shared objects. For concurrent modifications of the agenda, you should provide a SHARED EDITOR* for manipulating the agenda. The SHARED EDITOR* allows users to create new agenda items or edit existing agenda items. Changes to agenda items are instantly visible to all users of the shared editor. Users who are currently not using the shared editor or who currently have no focus on the modified item will receive CHANGE NOTIFICATIONS*. The SHARED EDITOR* can be connected with an EMBEDDED CHAT* so that users can discuss the content. An alternative for

asynchronous discussion of an agenda item is to add a threaded discussion to each agenda item object. Users can express their agreement with agenda items by means of an integrated (agenda item specific) VOTE*.

Drawbacks: If the agenda is created by the group, the meeting owner must still take care that it fulfills the requirements described, see NO AGENDA – NO MEETING and WHY SHOULD I BE THERE. The meeting owner should in this case act as a MODERATOR* to ensure that the meeting is still in line with the general meeting goals.

Related Patterns:

- WHY SHOULD I BE THERE: The discussion of the agenda can help the participants to better understand why they should be there.
- SHARED FILE REPOSITORY* and ROOM* can both be used to store the agenda and make it accessible.

Additional meeting patterns

Besides the IT'S MY AGENDA – IT'S MY MEETING pattern, other patterns were relevant for our interventions at *Alpha*. The complete meeting pattern language (Schümmer & Tandler, 2008) currently contains 21 patterns. For space reasons, we have only included one pattern in full length in this paper. The additional relevant patterns are shown in Figure 3 and summarized in the remainder of this section.

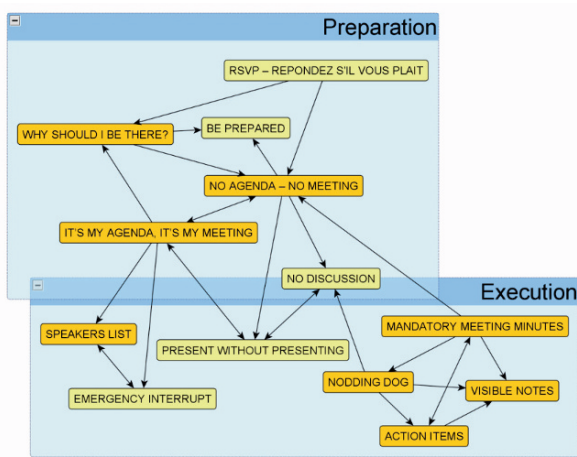


Figure 3: Excerpt of the meeting pattern language.

Figure 3 groups the patterns in two sections (Preparation, Execution). Patterns in the preparation cluster are mainly used before the meeting. They have the goal of helping the facilitator to create the most appropriate agenda and identify those people who can contribute to the agenda in order to make the meeting effective. The execution section contains patterns mainly applied during the meeting.

Obviously, results of patterns in the preparation cluster are used again during meeting execution (e.g., the meeting agenda). Dark colored patterns were selected as starting points for designing interventions in the *Alpha* case (see next section).

Patterns connected by an arrow are often used in combination. For instance, the ACTION ITEMS pattern is supported by the VISIBLE NOTES pattern: when participants establish a culture of creating action items for decisions, it is a good practice to show the action items during the meeting and make them part of the meeting minutes.

NO AGENDA – NO MEETING: Ensure that every meeting has an agenda. If not, cancel the meeting (motivation, clear goals, information, and focus).

WHY SHOULD I BE THERE: Give a reason why the receiver of an invitation should attend the meeting (motivation, clear goals, respect, and participation).

RSVP – RÉPONDEZ S'IL VOUS PLAÎT: Ask invited participants to respond to an invitation and modify the meeting plans if key participants cannot attend the meeting (reliability and participation).

BE PREPARED: Ensure that participants are able to prepare for a meeting (reliability, information, efficiency, and focus).

IT'S MY AGENDA, IT'S MY MEETING: Involve participants in the agenda creation (motivation, clear goals, and participation).

NO DISCUSSION: Plan agenda items without any discussions where participants only vote on a result (efficiency, focus, communication, and results).

PRESENT WITHOUT PRESENTING: Distribute presentation material before the meeting and thereby reduce the time spent on presentations in the meeting (information, efficiency, and focus).

SPEAKERS LIST: Manage a list of speakers and keep this list visible to all participants in order to facilitate awareness on upcoming contributions (respect, communication, and participation).

EMERGENCY INTERRUPT: Allow participants to signal a high priority speaker's request if their contribution addresses meta-issues, such as the end of the discussion or the request to vote for a final decision (efficiency, clear goals, focus, and results).

MANDATORY MEETING MINUTES: Ensure that decisions and action items are captured in minutes (information, clear goals, and results).

VISIBLE NOTES: Show the minutes to all participants while they are written and by that create a shared understanding of the discussion history (efficiency, focus, communication, and results).

ACTION ITEMS: Collect actions and responsibilities required for implementing a decision in order to ensure that these decisions are implemented after the meeting. Check the status of the action items periodically (reliability, clear goals, efficiency, focus, and results).

NODDING DOG: Quickly assess the level of commitment in the group so that the discussion can be stopped when agreement is reached (communication, participation, and results).

From indications to concrete interventions

The lack of properties detected in the ethnographic studies focused our interventions on five patterns, all related to the handling of issues and agendas. The selection of the patterns was driven by those properties that caused most problems in the observed meetings: the lack of *information* that would have been important for preparing for the meeting, differing levels of *participation* that caused participants to feel bored and consider the meeting as wasted time, and the lack of a shared artifact that documents the *results* of the meeting. In addition, we considered patterns that help to better *motivate* participants, mainly by means of improved information on the meeting's *goals*.

Three patterns were selected from the preparation cluster: The NO AGENDA – NO MEETING pattern was selected to convince the manager to provide more details on the topics of the meeting before the meeting takes place. In addition, the managers of the meetings were asked to explain the wanted participation of invited persons as stated in the WHY SHOULD I BE THERE pattern. The IT'S MY AGENDA, IT'S MY MEETING pattern was selected in order to give project members an opportunity to influence the agenda creation and bring up topics that were relevant for them. All three patterns have their focus on improving motivation, awareness of goals, and participation.

The selection of patterns for meeting execution was influenced by the fact that the VISIBLE NOTES pattern was already (partially) in place at *Alpha*. The meeting pattern language, however, suggests that notes taken should be accessible for all participants after the meeting (MANDATORY MEETING MINUTES). We thus proposed to fully implement the MANDATORY MEETING MINUTES pattern and to store the minutes in a shared space accessible to all project team members. The same advice applied to the handling of action items: again, we proposed to make ACTION ITEMS accessible to all team members and trace them through meetings.

The workshop at *Alpha*

In order to introduce the patterns, we conducted a two-day workshop with a selected project team (5 people) of *Alpha*. During the workshop, the participants were trained in detecting the need for the patterns as well as using wiki technology for supporting the social solutions of the patterns.

On day one, the participants discussed the patterns from the pattern language. We asked them to report on their experiences with meetings at *Alpha* and thereby helped them to map the abstract problem descriptions of the meeting patterns to the concrete context of *Alpha*. The discussion showed that the problems we observed were no single instances. Instead, all participants of the workshop reported that they have experienced the problems over and over again. Participants started to think about how the patterns could be used in their daily business utilizing the socio-technical infrastructure that was in place at *Alpha*. At

the end of day one, the participants were introduced to a wiki solution for the selected patterns.

The wiki solution based on the CURE wiki (Haake et al., 2005) models agendas, agenda-items, and action items as wiki pages. Special overview pages available in the CURE wiki were used to aggregate information from the individual agenda-item pages and the action item pages. Unlike most wikis, CURE supports in place editing of information visible in the overview pages.

The system was used in an artificial meeting setting on the second day. The goal was to experience the patterns in practice. Therefore, the project team members were confronted with a concrete task (“Organize the 20th anniversary event for *Alpha*”) and had to organize the meeting in which this task was addressed. One workshop participant took the role of the meeting facilitator and created a first draft of an agenda in a new meeting workspace. He then invited the other participants to the meeting space. CURE supports this by issuing virtual keys to these users (Haake et al., 2004).

Together with the invitation, the meeting organizer was able to provide a free text comment providing more details why participation is important (WHY SHOULD I BE THERE). This feature was not used as expected. Participants wrote messages like “please join our planning meeting” but did not consider the concrete roles of the invited participants. We cannot exactly name the reason for this but have two hypotheses: First of all, the observed setting was artificial. Roles were not clearly set up and it was thus difficult to identify concrete responsibilities for the participants. Secondly, the users reported that the usability of the invitation process could be improved. There is especially a missing link between the invitation and the agenda items referenced in the invitation.

After accepting the invitation, the participants could enter the virtual meeting space in the wiki and help to improve the agenda. We could observe that the participants brought in their ideas for agenda items that would be important for the meeting. This way, the agenda evolved until the participants were satisfied with the topics for the meeting.

Once the agenda was finalized, the team started the meeting after a short coffee break. Again, this is not the same setting as it would be present in a real meeting where participants have longer times between the finalization of the agenda and the meeting (in order to BE PREPARED). Additional findings were required to find out whether or not the preparation process was improved by the applied patterns. We conducted a survey at the end of the project asking also employees of *Alpha* for their perceived impact of the project. *Alpha* reported that setup time for meetings was reduced by two thirds. We cannot say to what extent our agenda handling patterns influenced this judgment but can consider this as positive feedback on the whole meeting preparation process.

During the meeting, the agenda items were used as a skeleton for MANDATORY MEETING MINUTES. The minutes were projected to a wall of the meeting room

(VISIBLE NOTES). Figure 4 provides an example of a meeting minutes page that was created during the workshop. Participants used a comment field present in each agenda item to summarize the discussions. In addition, they added ACTION ITEMS (on the bottom of Figure 4) to the minutes.

Minutes: Minutes
 These minutes capture decisions made in the meeting: [Kick off meeting](#)

Results from the individual agenda items

#	Name	Minutes
1	Decide on location	Björkhaga Hotell in Mulsjö has a nice conference room. We could also go to the Mulsjö missionskyrka.
2	Organize cake	The group prefers chociate based cakes.
4	Decision on music	

Global meeting notes

Action items



Name	Responsible	Due	Status	New	v
Book the band	 Gunnar Svensson	2007-10-11	running	YES	
Do more cookies	 Till Schümmer				
<input type="text" value="Check availability of locations"/>	...			YES	

Figure 4: Supporting the IT'S MY AGENDA - IT'S MY MEETING, VISIBLE NOTES, MANDATORY MEETING MINUTES, and ACTION ITEMS patterns in CURE.

After the workshop, the participants were asked for free-form feedback on the meeting patterns and the support technology. This feedback showed that the proposed social processes were understood by the participants and considered as valuable advice for improving their meeting culture. Some participants stated that they would start to modify parts of their behavior based on the given advice.

They were more skeptical with respect to the supportive technology, mainly for two reasons: Firstly, it is not integrated with the IT infrastructure of *Alpha* but a stand-alone wiki. A tighter integration would be needed if the technology support was to be deployed in the whole organization. Secondly, the participants detected several interaction breakdowns in the presented implementation. It is, e.g., difficult to trace the life cycle of an agenda or action item when it is discussed in more than one meeting. This was also true for the solution that was in place at *Alpha* before our intervention, but now that the participants became aware of this (also due to the ACTION ITEMS pattern that explicitly discussed this issue) they wanted to have a solution that allows follow-up activities for agenda or action items.

Another difficulty with the infrastructure was a lack of coordination mechanisms. Participants especially requested the presence of a SPEAKERS' LIST

and a mechanism for quick polls (NODDING DOG). These patterns address problems that we could observe especially in the distributed meetings during our ethnographic observations. They were not so critical in the co-located workshop setting. But the fact that the participants asked for such an integrated functionality is a further indication that the principles of the patterns were understood and the potentials for an improved meeting support system became clear.

LivingAgendas – an integrated meeting support system

The requests of the workshop participants led to the design of a prototype called *LivingAgendas*. It is a web-based meeting support system that supports groups in the collaborative creation of an agenda and in implementing the agenda during a real meeting. We put special attention on the interaction breakdowns that were detected with the previous wiki-based prototype. Figure 5 shows how the prototype was used in a real meeting. Before the screenshot was taken, the meeting organizer invited the participants. The system supports this by providing forms for explaining the need for participation to the invited person (not shown in Figure 5). All invited participants were allowed to propose changes and additions to the agenda. *LivingAgendas* enables automatically recalculating the time schedule and sending updated invitations to required participants.

During the meeting, the agenda is used as a template for the meeting minutes. Figure 5 shows the personal view of the participant named “Peter T.” who is currently taking the notes for the first agenda item. On the right part of Figure 5, one can see two lists that ease in-meeting coordination. The SPEAKERS LIST shows all participants that want to contribute to the current agenda item. Meeting members can add themselves to this list by pressing the “request to speak” link. They can also remove themselves from the list (by clicking on the “ok” link). The time information behind each participant is an estimate on how long it will take until the person gets the floor.

The second box on the right part of Figure 5 supports voting, both for testing the group’s current opinion (NODDING DOG) and making the final decision on a topic. Votes can be added to each agenda item by following the “Add Decision” link (also before the actual meeting takes place). As long as the facilitator does not forbid test votes, the voting is available immediately and interim results become instantly visible.

Once a decision is taken, the system suggests creating action items for implementing the decision. In Figure 5, the first agenda item has two associated actions. The actions are tracked automatically and added to the next status meeting for exchanging results or status information on the action items. This supports the entire life cycle of important action items or their corresponding agenda items. The same is true for tracking related agenda items over time. Both,

meeting preparation as well as meeting execution thus contribute to the meeting history of the team.

The screenshot displays the 'LivingAgendas' web application interface. At the top, there are navigation tabs for 'Preparation', 'Execution', and 'Follow Up'. The main content area is titled 'WP4 USE-CASE PLANNING' and includes the following details:

- Location:** Skype
- Date:** 2009-10-15 08:30:00 UTC
- Participants:** Jörg Haake, Hilda Telliöglu, Peter Tandler, Till Schümmer

On the left, there are sections for 'MY MEETINGS' (Paper Discussion, Pre-Project Meeting, Target Setting Meeting) and 'UPCOMING MEETINGS' (Supplier Conference, Division Meeting, Interview). On the right, there is a 'Speakers List' (Hilda Telliöglu (2*), Jörg Haake (3*), Till Schümmer (8*)) and 'Open Decisions' (Release of first prototype, Definition of test group).

The main agenda items are:

- 08:30-09:15: Release of first prototype**
 - Goals: - negotiate core use cases
 - Duration: 45
 - Facilitator: Till Schümmer
 - Notes: We discussed whether or not the main focus should be on collaborative editing. Selecting this scenario would allow us to improve tools for tighter collaboration. However, most of the collaboration up to now takes place in an asynchronous setting.
 - Decisions: -> Use collaborative Editing in Use-Case 1 (Yes: 2; No: 2; Don't care: 0)
 - Actions: -> Evaluate available shared editors that can be embedded in web-based systems. (Assigned: Hilda Telliöglu)
 - Create paper-based prototy (Assigned: Till Schümmer)
- 09:15-09:40: Definition of test group**
 - Goals: - availability of engineers -availability of sales representatives
 - Duration: 25
 - Facilitator: Jörg Haake

Figure 5: The prototype *LivingAgendas*.

LivingAgendas supports direct manipulation whenever possible. Agenda-items can, e.g., be moved by dragging them to the new position in the agenda or votes can be submitted by pressing on the circle. Technically, it is a web-based client-server application that uses Ruby on Rails as well as JavaScript frameworks that allow a rich usage experience.

Anecdotal experiences with the prototype make us think that *LivingAgendas* may be able to improve our concepts of pattern-guided meetings not only in the context of research but also in “real” work environments.

Conclusions

In this paper we presented an ethnographic study and showed how the absence of fundamental meeting properties we derived can be used to identify problems in meetings. The analysis of the properties aids a better understanding of areas for intervention.

We proposed a socio-technical meeting pattern language that allows a co-evolution of both the social processes required in a meeting and the supportive infrastructure. The patterns guided concrete interventions for opening up the process of agenda creation for all participants. The interventions were implemented by users themselves. They tailored the environment and thereby added more structure to their meeting infrastructure. However, users also detected interaction breakdowns that arose when they tried a wiki on their own. These problems together with findings from other research on meeting interaction helped us to better focus onto the development of a new meeting infrastructure called *LivingAgendas*. Our patterns are not tool-dependent; they are identified and studied in cooperative work settings. With *LivingAgendas* we aim at facilitating the whole meeting life cycle.

Our current and future work comprises three main research and development areas: (1) we are working on a public release of *LivingAgendas* in order to test it with a broader audience, (2) we plan to run ethnographic studies for the use of the *LivingAgendas* system looking deeper into the socio-technical effects of the meeting patterns, and (3) we are working on extending and improving the meeting pattern language, especially for addressing additional challenges of in-meeting interaction. The early feedback from the users encourages us to continue this path helping to make meetings more efficient.

Acknowledgments

The work presented in this paper was (partially) supported by the European Commission in the research project MAPPER, EU project number 016527. The ethnographic studies in MAPPER have been mainly carried out by Gian Marco Campagnolo, Gianni Jacucci, Hilda Tellioglu and Ina Wagner. We thank our colleagues for the excellent collection of rich data that allowed us to analyze real life meetings in this paper. Special thanks are due to Peter Tandler for co-authoring the pattern language with one of the authors of this paper. We like to thank the EuroPLOP community that helped to improve the patterns used in this paper. Finally, special thanks are due to all other members of the MAPPER project who made this research possible.

References

- Alexander, C. (1977): *The timeless way of building*, Oxford University Press, New York.
- Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I. and Angel, S. (1977): *A pattern language*. Oxford University Press, New York.
- Bicharra G. A. C., Kunz, J. and Fischer, M. (2004): 'Cutting to the chase: improving meeting effectiveness by focusing on the agenda', *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW'04*, Chicago, Illinois, USA, November 06-10, New York, NY, pp. 346-349.
- Byrne, J.; Brandt, R. and Port, O. (1993): 'The virtual corporation', *Business Week*, vol. 8, pp. 98-102.
- Carroll, J. and Farooq, U. (2007): 'Community-based learning: Design patterns and frameworks', *Proceedings of ECSCW 2007*, Paris, pp. 307-324.
- de Vreede, G.-J., Vogel, D., Kolfshoten, G. and J. Wien (2002): 'Fifteen Years of GSS in the Field: A Comparison Across Time and National Boundaries', *Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)*.
- Doyle, M. and Strauss, D. (1982): *How to make Meetings Work*, Jove Books, New York.
- Guy, E. S. (2005): '"...real, concrete facts about what works...": integrating evaluation and design through patterns', *Proceedings of GROUP'05*, ACM, New York.

- Haake, J. M., Haake, A., Schümmer, T., Bourimi, M. and Landgraf, B. (2004): 'End-user controlled group formation and access rights management in a shared workspace system', *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW '04*, Chicago, Illinois, USA, November 06-10, ACM, New York, NY, pp. 554-563.
- Haake, A., Lukosch, S. and Schümmer, T. (2005): 'Wiki-templates: adding structure support to wikis on demand', *Proceedings of the International Symposium on Wikis, WikiSym'05*, San Diego, California, October 16-18, ACM, New York, NY, pp. 41-51.
- Herlea et al. (2001): 'Secure Meeting Scheduling With Agents', in R. Steinmetz et al. (eds.): *Proceedings of IFIP CMS 2001, Communications and Multimedia Security Issues of the New Century*, May 21, Darmstadt, Germany, pp. 327-338.
- Herrmann, T., Hoffmann, M., Jahnke, I., Kienle, A., Kunau, G., Loser, K. and Menold, N. (2003): 'Concepts for usable patterns of groupware applications', *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*, Sanibel Island, Florida, USA, November 09-12, GROUP'03. ACM, New York, NY, pp. 349-358.
- Jacucci, G., Tellioglu, H. and Wagner I. (2006): *Report Kongsberg Automotive, MAPPER, Model-based Adaptive Product and Process Engineering, IST/NMP Project*, 016527.
- Jay, A. (1993): 'How to run a meeting', in R. Baecker (ed.): *Readings in Groupware and Computer-Supported Cooperative Work*, Morgan Kaufmann Publishers, pp. 130-144.
- Jennings et al. (1995): 'Agent-based Meeting Scheduling: A Design and Implementation', *Electronic Letters*, vol. 31, no 5, IEEE Press, pp. 350-352.
- Ju et al. (2008): 'Range: Exploring implicit interaction through electronic whiteboard design', *Proceedings of the ACM 2008 Conference on Computer Supported Cooperative Work, CSCW'08*.
- Kolfschoten, G., Appelman, J., Briggs, R. and de Vreede, G. J. (2004): 'Recurring Patterns of Facilitation Interventions in GSS Sessions', *Proceedings of the 37th Hawaii International Conference on System Sciences*, IEEE Press.
- Kelsey, D. and Plumb, P. (2004): *Great Meetings! Great Results*, Hanson Park Press, Inc.
- Matson, E. (1996): *The Seven Sins of Deadly Meetings*, Fast Company, April. http://www.fastcompany.com/online/02/meetings_Printer_Friendly.html.
- Microsoft SharePoint (2008): <http://office.microsoft.com/en-us/help/HA011154391033.aspx>, <http://office.microsoft.com/home/video.aspx?assetid=XT010944071033&width=884&height=540&startindex=0&CTT=11&Origin=HA010930841033>
- Nunamaker, J. F., Dennis, A. R., Valacich, J. S., Vogel, D. R. and George, J. F. (1991): 'Electronic meeting systems to support group work', *Communications of the ACM*, vol. 34, no 7, July, pp. 40-61.
- Parker, G. and Hoffman, R. (2006): *Meeting Excellence: 33 Tools to Lead Meetings That Get Results*, Jossey-Bass.
- Schümmer, T. and Haake, J. (2009): 'Shaping Collaborative Work with Proto-patterns', in V. Pipek, M. B. Rosson, B. de Ruyter and V. Wulf (eds.): *End-User Development. 2nd International Symposium, IS-EUD 2009*, Siegen, Germany, Springer, LNCS 5435, pp. 166-185.
- Schümmer, T. (2005): *A Pattern-Approach to End-User Centered Groupware Development*. Eul-Verlag, Lohmar, Cologne.
- Schümmer, T. & Lukosch, S. (2007): *Patterns for Computer-Mediated Interaction*. John Wiley & Sons., Chichester, UK.
- Schümmer, T. and Tandler, P. (2008): 'Patterns for Technology-Enhanced Meetings', in: L. Hvatum and T. Schümmer (eds.): *Proc. of EuroPLOP 2007*, UVK, Konstanz, Germany, pp. 97-119.
- SMART Board™ interactive whiteboards: <http://www2.smarttech.com/st/en-US/Products/SMART+Boards/>
- Streibel, B. J. (2003): *The Manager's Guide to Effective Meetings*, McGraw-Hill, New York.
- Streitz, N.A., Tandler, P., Müller-Tomfelde, C. and Konomi, S. (2001): 'Roomware: Towards the Next Generation of Human-Computer Interaction based on an Integrated Design of Real and Virtual Worlds', in J. Carroll (ed.): *Human-Computer Interaction in the New Millennium*, Addison-Wesley, pp. 553-578.
- Tandler, P. (2008): *Synchronous Collaboration in the Age of Ubiquitous Computing*, VDM Verlag Dr. Müller.
- Vivacqua et al. (2008): 'Assisting meeting facilitation through automated analysis of group dynamics', *The 12th International Conference on Computer Supported Cooperative Work in Design, CSCWD 2008*, pp. 951-956.
- Yankelovich, N., Walker, W., Roberts, P., Wessler, M., Kaplan, J. and Provino, J. (2004): 'Meeting central: making distributed meetings more effective', *Proceedings of CSCW'04*, ACM Press, New York, pp. 419-428.

Analyzing Multimodal Communication around a Shared Tabletop Display

Anne Marie Piper and James D. Hollan

Department of Cognitive Science, University of California, San Diego, USA
ampiper@hci.ucsd.edu, hollan@hci.ucsd.edu

Abstract. Communication between people is inherently multimodal. People employ speech, facial expressions, eye gaze, and gesture, among other facilities, to support communication and cooperative activity. Complexity of communication increases when a person is without a modality such as hearing, often resulting in dependence on another person or an assistive device to facilitate communication. This paper examines communication about medical topics through *Shared Speech Interface*, a multimodal tabletop display designed to assist communication between a hearing and deaf individual by converting speech-to-text and representing dialogue history on a shared interactive display surface. We compare communication mediated by a multimodal tabletop display and by a human sign language interpreter. Results indicate that the multimodal tabletop display (1) allows the deaf patient to watch the doctor when she is speaking, (2) encourages the doctor to exploit multimodal communication such as co-occurring gesture-speech, and (3) provides shared access to persistent, collaboratively produced representations of conversation. We also describe extensions of this communication technology, discuss how multimodal analysis techniques are useful in understanding the affects of multiuser multimodal tabletop systems, and briefly allude to the potential of applying computer vision techniques to assist analysis.

Introduction

Loss of hearing is a common problem that can result from a variety of factors (e.g., noise, aging, disease, and heredity). Approximately 28 million Americans have significant hearing loss, and of that group, almost six million are profoundly deaf (NIDCD, 2008). A primary form of communication within the United States deaf

community is American Sign Language (ASL). ASL interpreters play a central role in enabling face-to-face communication between deaf and hearing individuals. For the deaf population fluent in ASL, communicating through an interpreter is an optimal choice for many situations. Interpreters, however, are expensive and in many situations not available. Furthermore, though interpreters are bound by a confidentiality agreement, the presence of a third person in a private conversation may reduce a deaf person's comfort and inhibit their willingness to speak candidly. These factors are especially relevant for the topic of our current analysis: medical conversations between a deaf patient and a hearing, non-signing doctor.

We designed and evaluated Shared Speech Interface (SSI), a multimodal tabletop application that facilitates communication between a deaf and hearing individual. The application was designed to provide private and independent communication within the context of doctor-patient consultations. While our initial findings indicate that communicating through a multimodal tabletop display is both feasible and desirable for deaf individuals (Piper and Hollan, 2008), it is not yet clear how the tabletop display affects communication on a cognitive and social level. This paper presents a micro-analysis of interaction between deaf and hearing individuals to begin to address questions regarding communication, coordination, and cognition. Our analysis examines speech, gesture, eye gaze, and device interaction involving the doctor, patient, and sign language interpreter. We find that the digital table provides dialogue with properties that are not available in conversation through a human interpreter. Specifically, the digital table transforms ephemeral dialogue into a lasting form that allows the deaf individual to better attend to the speaker, supports co-occurring gesture-speech by the hearing user, and provides a shared visual record of conversation.

Deaf Communication

Deaf individuals living in a hearing world face communication challenges everyday and often rely on other people or devices to assist communication. While not all deaf or hearing impaired individuals use sign language, sources estimate that ASL is the fourth most widely used language in the United States (NIDCD, 2008). Sign language interpreters are a common solution for facilitating communication between deaf and hearing individuals, but access to an interpreter requires foresight and can be expensive. While interpreter services are important, they raise issues of privacy in communication. The Deaf community in many locations is small and well-connected. It is not uncommon for a deaf person to know the interpreter, which creates concern for very personal conversations. The interpreter scheduled on a given day may also be of the opposite gender, making discussion of certain medical issues even more uncomfortable. Face-to-face communication through an interpreter requires the deaf individual to focus their attention on the interpreter rather than the speaker. Taking notes during conversation involving an interpreter is also challenging because the deaf individual must pay close attention to the interpreter and cannot easily look down to make notes on paper. Not all deaf individuals

know how to read and write in a spoken language such as English, but those who are proficient may use hand written notes to communicate in the absence of an interpreter. Communication with the hearing world is further complicated because sign languages are not simply visual forms of spoken languages. Instead, each sign language has its own unique grammatical and syntactical structure, making a spoken language a second language for many deaf individuals.

Technology has transformed communication for the Deaf community. Telephone use was impossible for deaf individuals until the adaptation of the Teletype machine (TTY) which allowed individual lines of keyboard entry to be transmitted over phone lines. Adoption of the TTY, its subsequent electronic versions, and now the personal computer, made typing an essential mode of communication within the Deaf community. Researchers have developed a variety of technologies to address communication barriers between the deaf community and hearing world. As early as 1975, researchers began investigating how cooperative computing environments, such as early forms of instant messenger, could facilitate communication between deaf and hearing individuals (Turoff, 1975). More recently, human-computer interaction researchers have examined how mobile devices (e.g., Cavender et al., 2006), tablet computers (Miller et al., 2007), and browser based technologies (Schull, 2006) can augment communication for deaf individuals. While these solutions address various communication challenges for deaf individuals, none address face-to-face communication around a single shared display.

Multimodal Tabletop Displays

Digitally enhanced tabletop displays are growing in appeal and availability. The ability to receive multiple simultaneous touch inputs from a number of people makes tabletop displays a promising technology for facilitating face-to-face group interaction. Within the field of human-computer interaction, substantial attention is given to how tabletop displays can support face-to-face communication and mediate group social dynamics (see Morris, 2006, for a review). Compared to vertical displays such as a computer monitor or wall mounted display, tabletop displays result in more equitable interaction and shared responsibility by group members (Rogers and Lindley, 2004). Recently, there has been growing interest in multimodal multitouch tabletop systems. A multimodal tabletop system accepts touch along with speech and/or eye gaze as input to the system. Tse and his colleagues explored how multimodal tabletop systems support gaming, pair interaction around a multimodal tabletop display, and techniques to wrap single-user applications so they include multimodal interaction (2007). Researchers have examined a variety of tabletop group work issues with hearing populations, but until recently with the Shared Speech Interface project (Piper and Hollan, 2008), researchers had yet to examine tabletop computing scenarios with hearing impaired populations.

We developed Shared Speech Interface (SSI), a multimodal tabletop application that enables co-located face-to-face communication and cooperative activity between a hearing and deaf individual. The design of SSI exploits the affordances

of multimodal tabletop displays while addressing communication needs between a deaf patient and a hearing, non-signing medical doctor. Consultations with physicians often involve visuals such as medical records, charts, and scan images. Interactive tabletop displays are effective for presenting visual information to multiple people at once without necessarily designating one person as the owner of the visual. Taking notes while meeting with a physician is problematic for deaf individuals because it requires simultaneously attending to the doctor's facial expressions, the interpreter's visual representation of speech, and notes on paper. A multimodal tabletop display allows the doctor and patient to maintain face-to-face contact while viewing a shared, interactive representation of their conversation and other visual materials.

SSI runs on a MERL DiamondTouch table (Dietz and Leigh, 2001) and uses the DiamondSpin toolkit (Shen et al., 2004). The DiamondTouch table is a multiuser, multitouch top-projected tabletop display. People sit on conductive pads that enable the system to uniquely identify each user and where each user is touching the surface. SSI supports conversational input through standard keyboard entry and a headset microphone. The system is currently English based. Audio captured from the microphone is fed into a speech recognition engine, converted from speech-to-text, and then displayed on the tabletop interface. Currently, SSI works for two users communicating in a face-to-face setting. The hearing user speaks into the headset microphone and the deaf individual enters speech through a standard peripheral keyboard. As the two individuals communicate, their speech appears on the tabletop display in the form of moveable speech bubbles. See Piper and Hollan (2008) for a detailed description of the system design.



Figure 1. A medical doctor and a deaf patient communicate using Shared Speech Interface.

Analysis of Multimodal Human Interaction

While a tabletop display is considered multimodal when it has multiple modalities of input (i.e., touch and speech, or touch and eye tracking), interaction with other people around a tabletop display is inherently multimodal. In this paper we use

video analysis techniques to closely examine the interplay between speech, gesture, and eye gaze as well as interaction with the device. Video analysis is routinely used to understand activity within naturalistic settings (e.g., Heath, 1986), but some laboratory studies also include analysis of multimodal human interaction data (e.g., Bekker et al., 1995; Kraut et al., 2003; Kirk et al., 2005). From a methodological perspective, Kirk et al. (2005) note the importance of studying laboratory data in an “ethnographic fashion.” Furthermore, Hollan et al. (2000) argue more directly for an integrated approach to human-computer interaction research based on theories of distributed cognition and a combination of ethnographic and experimental techniques.

Gesture in Co-located and Remote Interaction

There is a growing interest in co-located gestural interaction and its relevance to the design of cooperative computing systems. Tang (1991) noted the pervasive nature of hand gestures in a group drawing activity and indicated the need to better understand this activity in relation to the people and artifacts in a co-located workspace. Bekker et al. (1995) studied gestures as a way of informing the design of cooperative systems. Kraut et al. (2003) examined how visual information, especially deictic reference, enabled situational awareness and conversational grounding in face-to-face, video-based, and audio-based interaction.

The horizontal form factor of tables has unique affordances for group work compared to vertically mounted displays. Work by Rogers and Lindley (2004) noted an increased use of gesture when groups interacted around a tabletop display compared to a whiteboard display. In another study, Rogers et al. (2004) found that touching a display with fingers has ancillary benefit for group work such as supporting turn-taking. With respect to gesture, Tse et al. (2007) provided similar observations of pairs interacting around a multimodal tabletop display. They noted that “speech and gesture commands serve double duty as both commands to the computer and as implicit communication to others.”

A number of systems examined how representing nonverbal behaviors such as gesture and eye gaze across remote environments affects interaction (e.g., Tang and Minneman, 1990, as an early example). Related to gesture analysis, Kirk et al. (2005) examined how specific hand gestures within the context of remote cooperative activity promote awareness and coordinate object focused actions. Similarly, Luff et al. (2006) examined how people working remotely use pointing gestures to coordinate and align themselves around objects of interest.

Gesture Analysis

The term *gesture* is polysemous for human-computer interaction researchers interested in touch-sensitive surfaces. On one hand, gestures are commands to a computer system administered by touching or moving an object, finger, or hand on an interactive surface. In a more traditional sense, the term gesture refers to the way

in which people move or use their body as a means of communication or expression with oneself or others. This section focuses on this latter meaning of gesture. Recently there has been a growing interest in using gesture analysis to understand communication between people (McNeill, 1992; Kendon and Muller, 2001) and within cooperative work environments (Goodwin and Goodwin, 1996; Hindmarsh and Heath, 2000; Zemel et al., 2008). This is largely driven by a theoretical shift from considering gesture as peripheral to human interaction to viewing gesture as central to communication and thought. Kendon (1980) was one of the first to articulate the perspective that speech and gesture are inextricably linked. McNeill proposed a theory that speech and gesture involve a single conceptual source (McNeill, 1985, 1992). He posits that speech and gesture acts develop together. This and related work (McNeill, 1992; Goldin-Meadow, 2003) provide a foundation for using speech and gesture as a way to understand cognitive activity. Furthermore, gesture can indicate underlying reasoning processes that a speaker may not be able to articulate (Goldin-Meadow, 2003), and thus a better understanding of gesture promises to play a crucial role in teaching and learning (see Roth, 2001, for a review).

For the purposes of our discussion and in agreement with practices of gesture researchers, we examine gesture as spontaneous movements of body or hands that are often produced in time with speech but may also occur in the absence of verbal utterances (see McNeill, 1992). Actions such as head scratching or moving an object in an environment are not considered gestures. In our analysis we pay particular attention to gestures that communicate and mediate activity. We classify gestures into David McNeill's widely accepted categories of beat, deictic, iconic, and metaphoric gesture (1992). Examining the frequency and patterns of various gesture types provides potential insight into how people exploit their bodies and environment to assist communication during multimodal tabletop interaction.

Within gesture research, sign language is considered a separate class of communication. Each sign language has a specific syntactical and grammatical structure, and specific gestural forms within a sign language take on linguistic meaning. Communicating through sign language, however, does not preclude the use of spontaneous gestures as described above. In fact, signers use the same proportion of meaningful gesture as speaking individuals use in verbal dialogue (Liddell and Metzger, 1998). There is growing evidence that people – both hearing and hearing impaired – attend to and interpret information in gestures (Goldin-Meadow, 2003; Cassell et al., 1999; Beattie and Shovelton, 1999).

Eye Gaze Analysis

In addition to gesture, other nonverbal interaction such as eye gaze can provide insight into communication. Early work by Kendon (1967) gives a history of gaze research and describes the function of gaze as “an act of perception by which one interactant can monitor the behavior of another, and as an expressive sign and regulatory signal by which he may influence the behavior of the other.” Change in gaze

direction such as looking away while speaking and then back to the listener at the end of an utterance gives listeners information about turn-taking (Duncan, 1972, 1974; Duncan and Fiske, 1977). Eye gaze is also used to demonstrate engagement (Goodwin, 2000, 1981) as well as indicate attention and show liking (Argyle and Cook, 1976; Kleinke, 1986) during face-to-face interaction. Eye gaze, accompanied with or without gesture, is also used in pointing acts (Kita, 2003).

When working with deaf populations, understanding patterns of eye gaze is especially important. Direction of gaze indicates whether or not an individual is attending to visual forms of speech. In conversation, a deaf individual reading sign will maintain relatively steady gaze towards the person signing (Baker and Padden, 1978; Siple, 1978). Eye contact with the signer is a signal that the signer has the floor, and shifting gaze away from the signer can indicate a turn request (Baker, 1977). In American Sign Language, the direction of gaze can also be used for deictic reference (Baker and Padden, 1978; Engberg-Pedersen, 2003), and monitoring gaze direction may provide insight into accompanying interaction. Signers tend to shift gaze from the face of their listener to their own hands when they want to call attention to gestures, and it is common for the signer to look back up at their listener to ensure that they too are looking at the gesture (Gullberg and Holmqvist, 2006). Work by Emmorey et al. (2008) found that people reading sign language do in fact follow gaze down to the hands when a signer looks at his or her hands. In summary, eye gaze is an important aspect of multimodal interaction and understanding it may lead to innovation in cooperative multimodal technology design.

Experimental Setup

Eight deaf adults (mean age=33, stdev=11.4, range=[22,52]; 3 males) and one medical doctor (age=28, female) participated in a laboratory study. All eight deaf participants were born deaf or became deaf before the age of one. Three participants identified English as their native language and five identified ASL. All participants were fluent in ASL and proficient at reading and writing in English. The medical doctor had prior experience treating deaf patients but does not know ASL. None of the participants had used a tabletop display prior to participating in this study.

Deaf participants were given sample medical issues (e.g., about routine vaccinations for travel abroad or advice on losing or gaining weight) to discuss with the doctor. Each deaf participant worked with the same doctor, which resembles the real-world scenario where one doctor has similar conversations with multiple patients throughout the day. The patient and doctor discussed a medical issue using either the multimodal tabletop system (digital table condition) or a professional American Sign Language interpreter (interpreter condition). Each discussion prompt had a corresponding medical visual that was preloaded into the tabletop system (e.g., a map for discussion about foreign travel). A paper version of the visual was provided for the interpreter condition. Medical professionals helped to ensure that the discussion prompts reflected authentic conversations that might occur in normal patient interaction but whose content did not require participants to

discuss information that might be too personal. Deaf participants experienced both the digital table and interpreter condition. The order of conditions and discussion prompts was randomized between subjects. Each session was video taped by two cameras from different angles to capture participants' interactions with each other and the digital table. All sessions were conducted around a DiamondTouch table to keep the environment consistent; the tabletop display was turned off for interpreter condition. Three researchers were present for the testing sessions and took notes. Each conversation with the doctor lasted from seven to nine minutes.

Our research team reviewed over two hours of video data, and together we transcribed and coded key segments of interaction. We were careful to select segments of activity that are representative of behavioral patterns. Video data were transcribed using notation techniques by Goodwin (2000) and McNeill (1992). Brackets surround speech that is co-timed with a gesture, and bold face speech indicates the stroke of the gesture. Transcriptions involving the interpreter indicate the interpreter's speech on behalf of the deaf individual and are not a transcription of sign language used.

Results

Initial findings indicate that Shared Speech Interface is a promising medium for facilitating medical conversations (see Piper and Hollan, 2008, for more details), but how does the multimodal tabletop display shape communication? To answer this question, analysis focuses on four areas of co-located interaction. First, we examine patterns of gaze by the deaf individual as a way to understand their attention during interaction. Second, we present an analysis of gesture by the doctor to identify differences in how she exploits multiple modes of communication depending on the communication medium. Then we discuss how the deaf individual monitors multiple modalities of communication with an emphasis on co-occurring gesture-speech by the doctor. Lastly, we describe how the tabletop display provides persistent, collaboratively produced representations that can aid discussion in cognitively valuable ways.

Use of Eye Gaze

Video data reveal distinctly different patterns of eye gaze by the deaf individual when conversation is mediated by an interpreter compared to the multimodal digital table. Eye gaze is a particularly critical channel of communication for deaf individuals, as conversation is purely visual. Examining eye gaze data allows us to infer where the deaf individual is attending during communication. Our results show that when an interpreter is involved in communication, the deaf individual focuses gaze on the interpreter and glances only momentarily at the doctor, as expected per Baker and Padden (1978) and Siple (1978). We found that deaf participants in our study looked at the interpreter when they were reading signs (i.e., "listening") as well as when they were signing (i.e., "speaking"). Consider the following excerpt

of conversation from the interpreter condition. In this interaction, the doctor fixes her gaze on the deaf patient; however, the deaf patient focuses primarily on the interpreter and makes limited eye contact with the doctor. In both conditions, the doctor maintains eye contact with the patient throughout the conversation and uses eye gaze and backchannel communication (e.g., head nodding in center frame of Figure 2) to demonstrate attention and agreement with the patient's speech.

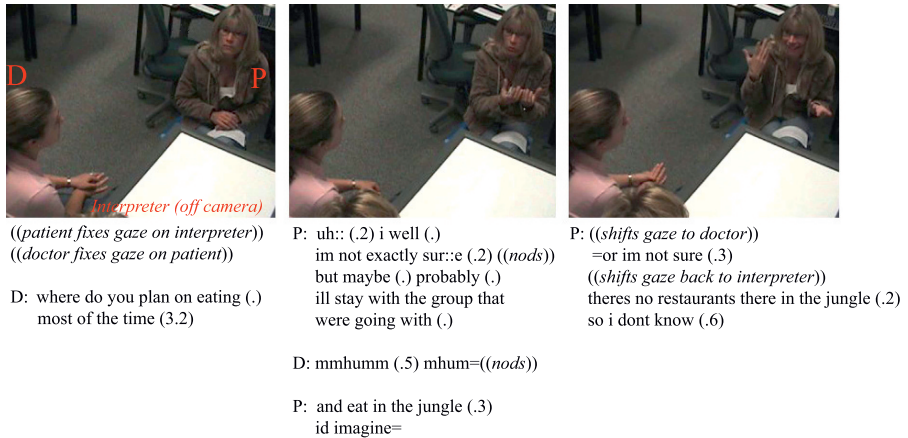


Figure 2. Doctor and patient communicating through interpreter. Patient watches interpreter while doctor looks at patient.

To elaborate this point, consider Figure 3 that illustrates the duration and patterns of eye gaze by this same individual. We highlight this case because the pattern illustrated here is typical for interaction. In the interpreter condition the patient fixes her gaze on the interpreter as needed for communication (Figure 3, grey areas in top bar graph). In contrast, communication through the digital table allows her to spend more time watching the doctor (Figure 3, black areas in bottom bar graph). As illustrated by Figure 3, when an interpreter mediates communication, this deaf patient makes quick one-second glances at the doctor and rarely holds gaze for longer than 3 seconds (gaze time on doctor: total=77sec, mean=2.1, stdev=2.0; gaze time on interpreter: total=293sec, mean=8.0, stdev=7.3). This is likely an attempt to demonstrate that she is attending to the doctor without signaling to the interpreter that she would like a turn to speak, as a sustained shift in eye gaze in sign language communication indicates a turn request (Baker, 1977). In the digital table condition, the patient makes frequent shifts in gaze between the doctor and tabletop and looks at the doctor for slightly longer intervals (gaze time on doctor: total=143sec, mean=3.0, stdev=2.6; gaze time on table: total=227sec, mean=4.9, stdev=7.7). The digital table requires the patient to look down for periods of time to type speech on the keyboard. Even with looking down at the keyboard, the doctor in our study noticed a difference in eye gaze by the patient. In a follow-up interview she said:

The physician patient interaction involves more than just words. Body language is integral to the medical interview and provides key details into the patient's condition

and level of understanding. The inclusion of the interpreter forced the deaf patients to make eye contact with her rather than me, not allowing me to gauge whether information or a question I asked was understood as well as more subtle insights into the patient's overall comfort level.




Figure 3. Duration and patterns of eye gaze by the deaf patient during the Interpreter and Digital Table conditions.

Use of Gesture

Communication through the digital table allows the patient to look at the doctor instead of requiring constant focus on the interpreter. Since speech appears in a permanent form on the tabletop display, the urgency of attending to the visual representation of talk is reduced. This allows both the doctor and patient to attend to and exploit multiple modalities of communication. Voice recognition capabilities free the doctor's hands and enable co-occurring gesture-speech in a way that traditional keyboard entry does not afford. Research on synchronized gesture-speech indicates that this activity is often co-expressive and non-redundant, therefore providing interactants with multiple forms of information (McNeill, 1992). Consider another example of interaction in Figures 4. Here, the doctor recommends hand washing techniques to the deaf patient by exploiting multiple modalities of communication including speech, gesture, and eye gaze. First, the patient looks at the doctor as she says "I would recommend." Then the doctor adds her speech to the display and continues "that you wash your hands." Both the doctor and patient look down at the display. Then the patient, likely to demonstrate understanding, holds up his hands and nods his head. The deaf patient's action is an iconic gestural response to the doctor's speech (McNeill, 1992). As he gestures, he shifts his gaze from the tabletop to his hands, likely to call the doctor's attention to his gesture (Gullberg and Holmqvist, 2006; Emmorey et al., 2008).

The patient then looks back at the doctor (Figure 4 middle row, left) as she formulates a recommendation for the patient. She makes a hand rubbing gesture as

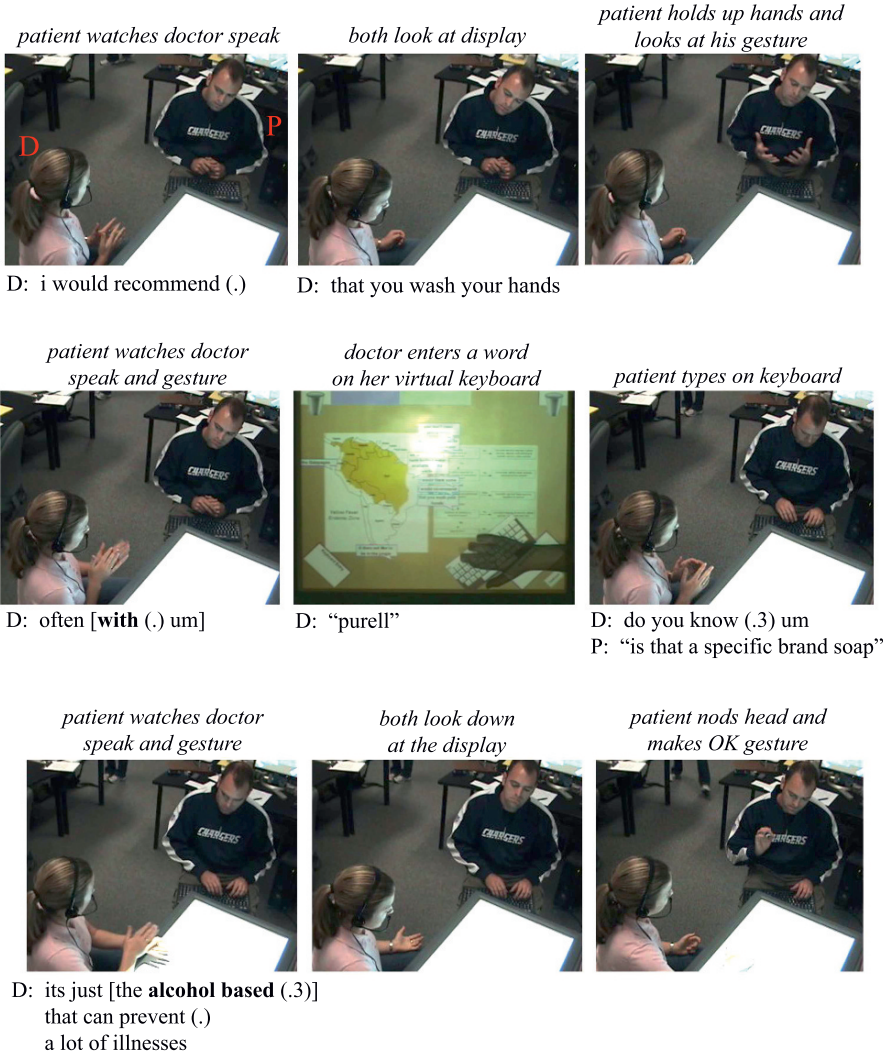


Figure 4. Doctor and patient communicate about hand washing through the digital table.

she says “with um.” Then she uses the virtual keyboard to type the word “purell.” The patient sees this word and responds by typing “Is that a specific brand soap?” His typing occurs simultaneously with the doctor’s speech (middle row, right frame of Figure 4). The doctor’s response (see Figure 4 bottom) demonstrates that she attends to the patient’s question for clarification. A critical moment in this interaction occurs in the bottom left image of Figure 4. The doctor and patient make eye contact as the doctor performs an iconic hand rubbing gesture timed with the words “alcohol based.” Her gesture communicates the method of use for hand san-

itizer, as alcohol-based sanitizers work by evaporating when rubbed into the hands. After this, both look down at the display to see the doctor's speech. Finally, the patient performs an emblematic "ok" gesture while nodding his head to show that he understands the doctor.


The doctor's carefully timed speech and gesture provide the patient with two pieces of information. First, her speech indicates the specific type of soap. Second, her gesture demonstrates how the soap is used. This information taken together yields a richer communicative form than either channel in isolation. This example demonstrates the importance of freeing the speaker's hands so that she is able to gesture as well as allowing the deaf individual to attend to the speaker's gestures instead of maintaining focus on the interpreter. In this example, and in others, we were struck by the highly coordinated use of speech, gesture, and eye gaze between the doctor and patient. The doctor's rich use of gesture to augment speech occurred often in interaction through the digital table. Similar use of gesture was *not* observed when the interpreter was present.

In a follow-up interview the doctor said that she intentionally tried not to gesture when the interpreter was present. She went on to explain that she did not want to compete with the interpreter for the patient's visual attention. In addition, interaction without the interpreter allowed the patient to frequently look at the doctor during communication, as is shown in Figure 3. This was a common pattern in the data. Having a larger percentage of the deaf patient's visual attention may have encouraged the doctor to elaborate her explanations with gesture (although this hypothesis needs to be examined with additional studies). Our analysis suggests that the multimodal tabletop system allows the doctor and patient to attend closely to each other's use of speech, gesture, and eye gaze as mechanisms for mediating communication. This also enables the doctor and patient to better monitor and exploit multiple modalities of communication such as co-occurring gesture-speech.

Monitoring Multiple Modalities of Communication

One challenge for deaf individuals involves monitoring multiple sources of visual information during conversation. Noticing and attending to co-occurring gesture-speech is a particularly challenging process when communication is mediated by an interpreter. Interpreter-mediated communication requires the deaf individual to notice co-occurring gesture-speech by the speaker and then put the speaker's gestures in context of the interpreter's gestural interpretation. Professionally trained interpreters are highly skilled, but they only occasionally replicate a speaker's gestures. Furthermore, through interviews with professional interpreters we found that their formal training does not specify when, if ever, they should replicate gestures made by the speaker. Overall, there were limited speech-gesture acts by the doctor in the interpreter condition, but this behavior did happen occasionally. Figure 5 is an example of the doctor using co-occurring gesture-speech. Here, she makes a fist like gesture (left) and then a two-handed iconic gesture (middle) to clarify portion size. Timing of speech and gesture is an issue, as the doctor completes each gesture

before the interpreter begins signing her speech. In this example, the interpreter did in fact recreate the doctor's gestures in context of her sign language interpretation but often the interpreter may not recreate the speaker's gesture, meaning that for at least a portion of communication the deaf individual must notice and attend to the speaker's gesture on their own. Even in cases in which the interpreter does recreate the gesture, it may not be formed or timed in exactly the same way as the original, thus creating interpretation challenges. In contrast, communication through the digital table provides opportunity for the deaf individual to look directly at the speaker's gestures, and as Figure 4 illustrates, gestures played an important role in establishing a shared understanding.



D: [cooked pasta should only be the size of your fist] (.4) D: [not the big bowls (.)] D: that are served in restaurants

Figure 5. Doctor uses gesture with her speech. Interpreter relays speech and gesture information..

Persistent, Collaboratively Produced Representations

Unlike other assistive technologies that enable communication between deaf and hearing individuals, the shared tabletop display provides a central focal point and space for establishing common ground (Clark and Brennan, 1991). The horizontal tabletop surface provides a space through which the doctor and patient cooperatively create, view, and manipulate representations of conversation. The shared conversation space allows the doctor and patient to gesture around and point to previous speech, thereby anchoring their gestures to objects (physical and virtual) in the environment (Clark, 2003). Referencing interface objects most often occurs through situated, context-specific pointing gestures (Goodwin, 2003). Both hearing and deaf participants used deixis to reference the material and symbolic world in front of them. With the interpreter, there is no record or explicit external representation of discourse. Consider Figure 6 (top row) where the doctor annotates a food pyramid diagram. Here, the doctor uses pointing gestures on a food pyramid diagram as she explains a balanced diet. The deaf patient must attend to both the interpreter's interpretation of speech as well as the doctor's pointing gesture occurring with her speech.

In this example, the doctor uses her speech and pointing gestures to walk the patient through parts of a food pyramid. Each time she points to a section of the

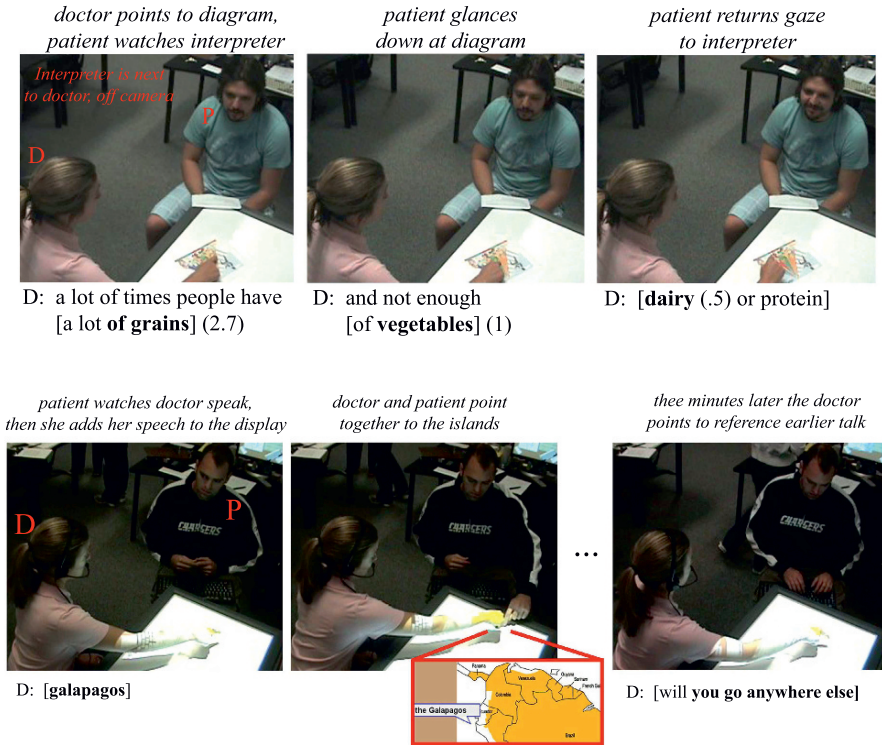


Figure 6. Top: Doctor points to parts of a diagram as she speaks. Patient monitors interpreter and doctor’s pointing gestures. Bottom: Using the digital table, the Doctor labels the Galapagos Islands on the map and then points to the speech bubble three minutes later.

diagram, she shifts her gaze to the table, likely an attempt to draw her listener’s attention to the diagram. Several minutes later the doctor references this diagram again to summarize her recommendation about a well-balanced diet, but the conversation and gestures she made to the patient are now only a memory.

The digital table stores collaboratively created representations of speech and allows users to rearrange and manipulate speech bubbles. Images in the top row of Figure 6 illustrate challenges with pointing to parts of a diagram while speaking; the digital table uniquely supports this form of interaction. We observed an interesting form of pointing that occurred through the strategic placement of speech bubbles. The tangible and persistent nature of speech bubbles affords certain interactions by serving as manipulatable cognitive artifacts (Hutchins, 1995). A speech bubble gains meaning beyond its literal text depending on how it is situated, or anchored, with respect to other parts of the activity. The canonical shape of speech bubbles, specifically the tail, allows the doctor and patient to use the objects as a pointing mechanism. That is, participants strategically placed speech bubbles around the display so that the tail of the speech bubble touched a relevant part of the background

or another speech bubble. Figure 6 (bottom center frame) provides an example of this behavior. In this interaction the doctor uses a speech bubble to label and reference part of a map. The patient mentions that he is traveling to the Galapagos Islands. The doctor says “Galapagos” as she points, and the patient points along with her to clarify the location. Subsequently, the doctor moves the “Galapagos” speech bubble to label the islands on the map. Then she uses this action to show that the islands are outside the Yellow Fever endemic zone (bottom center frame of Figure 6) and explain that the patient will not need the Yellow Fever vaccine. Conversation continues, and the topic changes. Approximately three minutes later the doctor comes back to “the Galapagos” speech bubble. She points to the speech bubble while asking, “will you go anywhere else?”

The persistent nature of speech along with the shared context of the tabletop display affords referencing both new and previously created external representations of speech. The persistent nature of speech also allows participants to review their entire conversation. Both the doctor and patients looked back over their previous conversation at some point during the activity. In a post-session interview, the doctor said, “It was good to look back at what I had covered with that particular patient,” and explained that, “[The digital table] would be helpful because it is not uncommon in medicine to have very similar conversations with different patients throughout the day.”

Discussion

Our analysis highlights differences in interaction between a deaf and hearing individual when communication is mediated by a multimodal tabletop display as compared to a human sign language interpreter. These differences reveal several trade-offs. Although speech recognition technology can not yet provide the richness and accuracy associated with translation by a competent interpreter, it does allow the doctor to exploit gesture for communicative purposes without fearing that she might distract the deaf individual from the interpreter. One example is the hand washing iconic display coinciding with speech depicted in Figure 4. In addition, transcribed speech-to-text allows the doctor and patient to have a shared record of conversation. This provides new artifacts in the environment, enabling pointing and other gestures (Roth and Lawless, 2002). Removing the time-critical visual demands of interpreter-mediated communication allows the deaf individual to focus more on the doctor while she is speaking. In turn, this helps the patient attend to the doctor’s speech-gesture acts and enables the doctor to better gauge patient understanding through increased eye contact. Speed of communication is another important trade-off issue. Current speech recognition is no match for a skilled interpreter. When using the speech recognition system, the doctor must speak slowly and carefully in order to help ensure accurate recognition. Time is also taken in selecting the appropriate alternative from the output of the recognition system and in correcting it when required. But necessitating slower dialogue on the part of the doctor is not an entirely negative outcome. Considering that English is a second language for many

deaf individuals, slowing the doctor's speech could in fact be a positive cognitive consequence of communicating through the tabletop display.

The SSI system technology has the potential to benefit multiple user groups and enable new cooperative computing applications. Shared displays, especially tabletop displays, are beneficial for a variety of group work tasks. Since inception of our project and the idea to visually represent conversation on a tabletop display, members of the Deaf community have mentioned numerous contexts in which this could be useful. Of these, the most frequently identified are counseling or therapy sessions, banking and financial services, meetings with an architect or interior designer, ASL and lip reading education, classroom group work, and even retail environments. Beyond the Deaf community, the cognitive affordances of SSI have implications for individuals with moderate hearing loss as well as unimpaired hearing users. The challenge of medical conversations is certainly not restricted to the Deaf community. Because of associated stress and other factors, it is easy to forget details to tell the doctor and even easier to forget specific instructions given during consultation. The affordances of SSI such as preloading questions for the doctor and referencing a transcript of a previous conversation extend to all populations. Similarly the ability to archive and subsequently revisit past multimodal conversations and collaborations has interesting potential to augment interaction.

The concepts behind SSI also have specific implications for user populations with other language-related communication barriers. For example, representing speech on a shared display has pedagogical benefits for language learning. Consider a case in which speech bubbles store textual and auditory information from a native speaking teacher and a student learning a second language. Here, both textual and auditory representations can be accessed in a shared collaborative context. The availability of visual and spatial representations of language also stand to benefit individuals with linguistic processing disabilities such as Aphasia or Apraxia. Language could take on a variety of representations including textual, auditory, and pictorial forms. For these individuals and other populations, a shared, co-located workspace has considerable promise to help in establishing common ground and assisting communication.

Conclusions and Future Work

Analysis of multimodal human interaction data is primarily used in ethnographic approaches to understanding everyday activity (e.g., Goodwin, 1981; Heath, 1986), but there is a growing interest in using multimodal analysis to understand the role of gesture occurring in experimental cooperative work settings (Bekker et al., 1995; Kraut et al., 2003; Kirk et al., 2005). We suggest that multimodal analysis can aid laboratory evaluations of tabletop technology as well as other cooperative work technologies in the following ways: (1) analysis of eye gaze provides a metric for understanding how people coordinate visual attention, (2) evaluation of gesture types and frequency of use provides a way to measure differences in interaction between experimental conditions, and (3) the interplay between speech, gesture, and

eye gaze can reveal cognitive and social consequences of new interactive media that would be difficult to detect with other methods.

Multimodal analysis, however, is tedious and extremely time-consuming. When analysis is so difficult, few analyses can be done and datasets are severely underutilized. Researchers come to have a large investment in the chosen data segments. Since each analysis may appear as an isolated case study, it can be difficult to know how common the observed phenomena may be. Larger patterns and contradictory cases can easily go unnoticed. Well-known human confirmation biases can affect the quality of the science when each analysis requires so much effort. The analyses presented in this paper, for example, resulted from a year-long iterative process of analysis of video and audio data to understand how differing communication media shapes interaction. This form of detailed analysis plays an increasingly central role in our ongoing investigation of tabletop display systems. One way our research group is addressing the difficulties of such analysis is by exploring techniques to assist video analysis. Applying computer vision techniques make it possible to tag video frames with certain characteristics of interest such as movement of hands or arms. We are currently evaluating computer vision methods for object recognition, face detection, and head pose estimation. For example, SIFT (Scale Invariant Feature Transform) (Low, 2004) is one popular and useful technique we are exploring. We see tremendous potential for computer vision techniques to assist video analysis for the types of data we report here and are exploring this as part of our ongoing work.

In this paper we have examined communication about medical topics through *Shared Speech Interface*, a multimodal tabletop display designed to assist communication between a hearing and deaf individual by converting speech-to-text and representing dialogue history on a shared interactive display surface. Specifically, we compared communication mediated by a multimodal tabletop display and by a human sign language interpreter. Results indicate that the multimodal tabletop display (1) allows the deaf patient to watch the doctor when she is speaking, (2) encourages the doctor to exploit multimodal communication such as co-occurring gesture-speech, and (3) provides shared access to persistent, collaboratively produced representations of conversation. Finally, we discuss extensions of our system and practical aspects of conducting a multimodal analysis of tabletop interaction.

Acknowledgments

Research is supported by a NSF Graduate Research Fellowship, NSF Grant 0729013, and a Chancellor's Interdisciplinary Grant. We thank our study participants, faculty and staff from UCSD Medical School, Whitney Friedman, and MERL for donating a DiamondTouch table.

References

Argyle, M. and M. Cook (1976): *Gaze and mutual gaze*. Cambridge University Press.

- Baker, C. (1977): 'Regulators and turn-taking in American Sign Language discourse'. In: *On the other hand: New perspectives on American Sign Language*. New York, pp. 215–236, Academic Press.
- Baker, C. and C. Padden (1978): 'Focusing on the nonmanual components of American Sign Language'. In: *Understanding Language through Sign Language Research*. New York, pp. 27–57, Academic Press.
- Beattie, G. and H. Shovelton (1999): 'Do iconic hand gestures really contribute anything to the semantic information conveyed by speech? An experimental investigation'. *Semiotica*, vol. 123, no. 1-2, pp. 1.
- Bekker, M. M., J. S. Olson, and G. M. Olson (1995): 'Analysis of gestures in face-to-face design teams provides guidance for how to use groupware in design'. In: *Proceedings of conference on Designing Interactive Systems (DIS)*. pp. 157–166.
- Cassell, J., D. McNeill, and K.-E. McCullough (1999): 'Speech-gesture mismatches: Evidence for one underlying representation of linguistic and nonlinguistic information'. *Pragmatics cognition*, vol. 7, pp. 1.
- Cavender, A., R. E. Ladner, and E. A. Riskin (2006): 'MobileASL: intelligibility of sign language video as constrained by mobile phone technology'. In: *Proceedings of conference on Computers and Accessibility (ASSETS)*. pp. 71–78.
- Clark, H. (2003): 'Pointing and Placing'. In: S. Kita (ed.): *Pointing: Where Language, Culture, and Cognition Meet*. Mahwah, NJ, pp. 243–268, Lawrence Erlbaum Associates.
- Clark, H. and S. Brennan (1991): 'Grounding in Communication'. In: L. Resnick, J. Levine, and S. Teasley (eds.): *Perspectives on Socially Shared Cognition*. Washington, APA Books.
- Dietz, P. and D. Leigh (2001): 'DiamondTouch: a multi-user touch technology'. In: *Proceedings of symposium on User Interface Software and Technology (UIST)*. pp. 219–226.
- Duncan, S. (1972): 'Some signals and rules for taking turns in conversations'. *Journal of personality and social psychology*, vol. 23, no. 2, pp. 283.
- Duncan, S. (1974): 'On the Structure of Speaker-Auditor Interaction During Speaking Turns.'. *Language in society*, vol. 3, no. 2, pp. 161.
- Duncan, S. and D. W. Fiske (1977): *Face-to-face interaction: Research, methods and theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Emmorey, K., R. Thompson, and R. Colvin (2008): 'Eye Gaze During Comprehension of American Sign Language by Native and Beginning Signers'. *Journal of Deaf Studies and Deaf Education*.
- Engberg-Pedersen, E. (2003): 'From Pointing to Reference and Predication: Pointing Signs, Eye-gaze, and Head and Body Orientation in Danish Sign Language'. In: S. Kita (ed.): *Pointing: Where Language, Culture, and Cognition Meet*. Lawrence Erlbaum Associates.
- Goldin-Meadow, S. (2003): *Hearing Gesture: How our Hands Help Us Think*. Harvard University Press.
- Goodwin, C. (1981): *Conversational Organization: Interaction Between Speakers and Hearers*. New York: Academic Press.
- Goodwin, C. (2000): 'Practices of Seeing, Visual Analysis: An Ethnomethodological Approach'. In: *Handbook of Visual Analysis*. London, pp. 157–182, Sage.

- Goodwin, C. (2003): 'Pointing as Situated Practice'. In: S. Kita (ed.): *Pointing: Where Language, Culture, and Cognition Meet*. Lawrence Erlbaum Associates.
- Goodwin, C. and M. Goodwin (1996): 'Formulating Planes: Seeing as Situated Activity'. In: *Cognition and Communication at Work*. pp. 61–95, Cambridge University Press.
- Gullberg, M. and K. Holmqvist (2006): 'What speakers do and what addressees look at: Visual attention to gestures in human interaction live and on video'. *Pragmatics cognition*, vol. 14, no. 1, pp. 53.
- Heath, C. (1986): *Body movement and speech in medical interaction*. Cambridge University Press.
- Hindmarsh, J. and C. Heath (2000): 'Embodied reference: A study of deixis in workplace interaction'. *Journal of Pragmatics*, vol. 32, no. 12, pp. 1855.
- Hollan, J. D., E. Hutchins, and D. Kirsh (2000): 'Distributed cognition: toward a new foundation for human-computer interaction research'. *ACM transactions on computer-human interaction*, vol. 7, no. 2, pp. 174–196.
- Hutchins, E. (1995): *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kendon, A. (1967): 'Some functions of gaze-direction in social interaction.'. *Acta Psychologica*, vol. 26, no. 1, pp. 22–63.
- Kendon, A. (1980): 'Gesticulation and Speech: Two Aspects of the Process of Utterance'. In: *The Relationship of Verbal and Nonverbal Communication*. p. 388, Walter de Gruyter.
- Kendon, A. and C. Muller (2001): 'Introducing: GESTURE'. *Gesture*, vol. 1, no. 1, pp. 1.
- Kirk, D., A. Crabtree, and T. Rodden (2005): 'Ways of the hands'. In: *Proceedings of European Conference on Computer Supported Cooperative Work (ECSCW)*. pp. 1–21.
- Kita, S. (2003): *Pointing: where language, culture, and cognition meet*. Lawrence Erlbaum Associates.
- Kleinke, C. (1986): 'Gaze and eye contact: A research review'. *Psychological Bulletin*, vol. 100, no. 1, pp. 78–100.
- Kraut, R. E., S. R. Fussell, and J. Siegel (2003): 'Visual information as a conversational resource in collaborative physical tasks'. *Hum.-Comput. Interact.*, vol. 18, no. 1, pp. 13–49.
- Liddell, S. K. and M. Metzger (1998): 'Gesture in sign language discourse'. *Journal of Pragmatics*, vol. 30, no. 6, pp. 657 – 697.
- Low, D. G. (2004): 'Distinctive image features from scale-invariant keypoints'. *International Journal of Computer Vision*, vol. 60, pp. 91–110.
- Luff, P., C. Heath, H. Kuzuoka, K. Yamazaki, and J. Yamashita (2006): 'Handling documents and discriminating objects in hybrid spaces'. In: *Proceedings of the conference on Human Factors in Computing Systems (CHI)*. pp. 561–570.
- McNeill, D. (1985): 'So You Think Gestures Are Nonverbal?'. *Psychological review*, vol. 92, no. 3, pp. 350.
- McNeill, D. (1992): *Hand and Mind: What Gestures Reveal about Thought*. University of Chicago Press.

- Miller, D., K. Gyllstrom, D. Stotts, and J. Culp (2007): 'Semi-transparent video interfaces to assist deaf persons in meetings'. In: *ACM-SE 45: Proceedings of the 45th annual southeast regional conference*. New York, NY, USA, pp. 501–506, ACM.
- Morris, M. R. (2006): 'Supporting Effective Interaction with Tabletop Groupware'. Ph.D. thesis, Stanford University, Stanford, CA.
- NIDCD (2008): 'National Institute on Deafness and Other Communication Disorders'. <http://www.nidcd.nih.gov>.
- Piper, A. M. and J. D. Hollan (2008): 'Supporting medical conversations between deaf and hearing individuals with tabletop displays'. In: *Proceedings of the conference on Computer-Supported Cooperative Work (CSCW)*. pp. 147–156.
- Rogers, Y., W. Hazlewood, E. Blevis, and Y.-K. Lim (2004): 'Finger talk: collaborative decision-making using talk and fingertip interaction around a tabletop display'. In: *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*. New York, NY, USA, pp. 1271–1274, ACM.
- Rogers, Y. and S. Lindley (2004): 'Collaborating around vertical and horizontal large interactive displays: which way is best?'. *Interacting with Computers*, vol. 16, no. 6, pp. 1133 – 1152.
- Roth, W.-M. (2001): 'Gestures: Their Role in Teaching and Learning'. *Review of Educational Research*, vol. 71, no. 3, pp. 365–392.
- Roth, W.-M. and D. V. Lawless (2002): 'When up is down and down is up: Body orientation, proximity, and gestures as resources'. *Language in Society*, vol. 31, no. 01, pp. 1–28.
- Schull, J. (2006): 'An extensible, scalable browser-based architecture for synchronous and asynchronous communication and collaboration systems for deaf and hearing individuals'. In: *Proceedings of the conference on Computers and Accessibility (ASSETS)*. pp. 285–286.
- Shen, C., F. D. Vernier, C. Forlines, and M. Ringel (2004): 'DiamondSpin: an extensible toolkit for around-the-table interaction'. In: *Proceedings of the conference on Human Factors in Computing Systems*. pp. 167–174.
- Siple, P. (1978): 'Visual Constraints for Sign Language Communication'. *Sign Language Studies*.
- Tang, J. C. (1991): 'Findings from observational studies of collaborative work'. *International Journal of Man-Machine Studies*, vol. 34, no. 2, pp. 143 – 160. Special Issue: Computer-supported Cooperative Work and Groupware. Part 1.
- Tang, J. C. and S. L. Minneman (1990): 'VideoDraw: a video interface for collaborative drawing'. In: *Proceedings of the conference on Human Factors in Computing Systems (CHI)*. pp. 313–320.
- Tse, E. (2007): 'Multimodal Co-Located Interaction'. Ph.D. thesis, The University of Calgary, Calgary, Alberta, Canada.
- Tse, E., C. Shen, S. Greenberg, and C. Forlines (2007): 'How pairs interact over a multimodal digital table'. In: *Proceedings of the conference on Human Factors in Computing Systems (CHI)*. pp. 215–218.
- Turoff, M. (1975): 'Computerized conferencing for the deaf and handicapped'. *SIGCAPH Comput. Phys. Handicap.*, no. 16, pp. 4–11.
- Zemel, A., T. Koschmann, C. Lebaron, and P. Feltovich (2008): "'What are We Missing?' Usability's Indexical Ground'. *Computer Supported Cooperative Work*, vol. 17, no. 1, pp. 63–85.

Status on Display: a Field Trial of Nomatic*Viz

Xianghua Ding, Donald J. Patterson
University of California, Irvine
[dingx,djp3]@ics.uci.edu

Abstract.

The use of personal status messages is becoming a part of popular culture through wide-spread instant messaging (IM) adoption, the growth of social networking websites and the increased connectivity provided by mobile phones. However, the implications of status broadcasting and people's behavior in the milieu of social life is still poorly understood. In this paper, we present the results of a field trial in which we examined how community members come to understand and appropriate a status broadcasting service into their daily use. We designed Nomatic*Viz, a situated large display showing people's location and status messages to complement an existing status message distribution tool called Nomatic*IM. Through a five month field study of its use we uncovered not only how it supports lightweight awareness of the community, but also how it participates in creating new spatial experiences and how people perform and negotiate self-representations through multiple simultaneous displays of personal status.

Introduction

As early as 1971, networked Unix computer users were using utilities to see the *status* of other mainframe users. By combining commands such as “who”, “ps”, and “finger” people had a way to describe their current workload, view other users' activities and account for the usage of shared computing resources. Gradually, a related idea was conceptualized, *presence*, to describe an indication that someone is in a digital or physical space (Fitzpatrick, 2003). This concept came about largely as the result of the increased availability of cameras, and other sensing peripherals.

Projects such as MediaSpaces (Bly et al., 1993) and Portholes (Dourish and Bly, 1992) both showed the value of this kind of awareness amongst remote collaborators.

As the internet has become more central to daily life, and as mobile computing has increasingly reduced the time when people are offline, these two concepts have become more intertwined (Patterson et al., 2008). Their hybridization has been seen in widespread distribution of personal status messages in popular culture, which is supported in social networking services such as Facebook¹, and Twitter². While these services provide support for quick status message authoring and sharing, other services support cross-service aggregation and redistribution such as FriendFeed³ and Ping.fm⁴.

Despite the widespread use of status broadcasting services in popular culture, the behavior and implications of its use in the milieu of social group life is still poorly understood. Smale and Greenberg noticed that users adapt their ID field on IM to broadcast information to their buddy list, and through examining these messages, identified different thematic categories (Smale and Greenberg, 2005). However, more work needs to be done to gain in-depth understanding of how users experience and practice status messaging. At the same time, researchers have studied automatically reporting sensed data including raw video, audio (Bly et al., 1993; Dourish and Bly, 1992) and a variety of sensed contextual data such as location (Brown et al., 2007), activity (Rowan and Mynatt, 2005), and IM availability (Terrell and McCrickard, 2006; Guzman et al., 2004) for their ability to support group awareness in a variety of settings, and have shown their value for improving coordination and connections. What happens, then, as status messages are made available not just for sharing with one's own buddies, but are also employed for awareness of a community in a shared space?

In this paper we report the results from a long term field trial of a distinctive status broadcasting system called Nomatic which repurposes status broadcasting to enhance awareness for a community. Figure 1 shows a usage model of the system. The black Nomatic*Viz display on the right, shows status messages and sensor data, such as location, activity and mood information. It displays the information in an ambiguous, anonymous and abstract manner in hopes of giving insights into community behavior while preserving privacy and promoting active engagement through data interpretation. The source for the status information is a context-aware software tool called Nomatic*IM which is installed on personal computers (usually laptops). Nomatic*IM uses a wide variety of built-in sensors (e.g., WiFi access points, ambient light, battery charging status) to categorize a user's context, to remind users to update their status when their context changes and to suggest status messages. The combination of these techniques are designed to support keeping status messages up-to-date with minimal cognitive burden.

1 www.facebook.com
 2 www.twitter.com
 3 www.friendfeed.com
 4 www.ping.fm

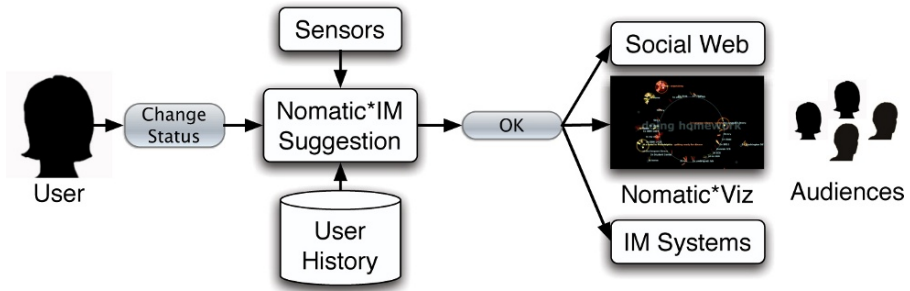


Figure 1. When a user wants to change status, possibly at the prompting of the system, Nomatic uses sensor data to provide suggested status messages. When a user selects a status message, it is sent to many different status broadcasting systems and Nomatic*Viz.

What displaying status messages in a community helps us to uncover is not just how it supports lightweight communal awareness, but more importantly, how the status display participates in producing new spatial experiences and how people perform with the display with simple text editing. We argue that the performative and performance aspects of awareness technologies, while often overlooked, play important roles in shaping people’s experiences with status broadcasting systems and other awareness technologies.

In the sections to follow, we discuss the relevant literature, the design, the deployment and the evaluation methodology of Nomatic*Viz. We then go on to synthesize a series of themes that emerged out of our qualitative analysis of the data. While these issues are specific to our particular case, they suggest social patterns that will help to frame the usage of other status broadcasting and awareness systems.

Related Work

Remote awareness and remote presence have long been important topics in the CSCW and related communities. They are usually viewed as mechanisms that improve collaborations and connections, especially for distributed groups. As mentioned previously, continuous streaming of video and audio was shown to foster ambient awareness in the Media Spaces and Porthole projects (Bly et al., 1993; Dourish and Bly, 1992). Later research on awareness started employing IM availability as an information source for ambient displays, and attempted to extend the availability information off the desktop and into our everyday environments (Guzman et al., 2004; Terrell and McCrickard, 2006).

What these different systems often share is the automatic capture and collection of information to minimize user distraction. By automatically capturing or rebroadcasting otherwise freely-available information, these systems effectively reduce the overhead that would be required if users had to keep the information up-to-date manually. Similarly, work on automatic sensing combined with context model-

ing technologies explores automatic ways to go beyond just reporting raw sensor data and instead assists in choosing an appropriate time to interrupt (Hudson et al., 2003), or choosing when to have a face-to-face meeting (Moran et al., 1999). The negative side of automatic broadcasting is the potential for inadvertent disclosure of information, particularly as the size of the audience grows larger (Patil and Lai, 2005).

Recent work has combined sensing with manual labeling to provide light weight automatic interpretations of sensor data for applications that are not predetermined. For example, the Awarenex project examined the user interface of a mobile status sharing system that associated manually entered labels with device detected locations (Tang et al., 2001). Reno allowed users to associate labels to cell tower connections and then activate rules based on entering those zones (Smith et al., 2005). Using a very different user interface, but similar underlying technology, the Whereabouts Clock allows users to associate three labels (work, home, school) with cell phone towers for the purpose of communicating vague locations to family members (Brown et al., 2007).

Nomatic*IM is like the latter systems in that it attempts to map labels to sensor data and to communicate those labels on behalf of users (Patterson et al., 2006). But unlike these systems it treats labels less strictly as a single description of location and more flexibly as multiple descriptions of context composed of “places”, “activities” and “moods”. The difference lies in relaxing the one-to-one mapping of physical locations to labels so that a location can be a “classroom” in the morning and a “conference room” in the afternoon. Similarly you might “be eating” in the kitchen at lunch, but “doing homework” late at night. It employs sensing to *automatically* label a current context, and also detects changes in context which acts as a reminder for status updates. For this field trial, Nomatic*IM was adapted to report sensor *and* status message information to the Nomatic*Viz large display in addition to publishing *just* status messages to a variety of IM systems, Twitter and Facebook.

Because Nomatic*Viz is a large display, recent investigations into how large displays enhance interactive and collaborative experiences are also germane. Work by Huang and Mynatt (2003), Churchill et al. (2003), Greenberg and Rounding (2001) and McDonald et al. (2008) have demonstrated that large displays have unique affordances. Compared to laptop or mobile phone displays, they are physically persistent, situated, and often shared, and as a result support low overhead awareness and information exchange, creating opportunities for conversations and community engagement and are usually the subject of many short “glances” rather than prolonged interactions.

About the Nomatic System

The whole Nomatic system is composed of two components, a context-aware software tool called Nomatic*IM which is installed on individual laptops and Nomatic*Viz, a display of participant’s contextual and status information located in

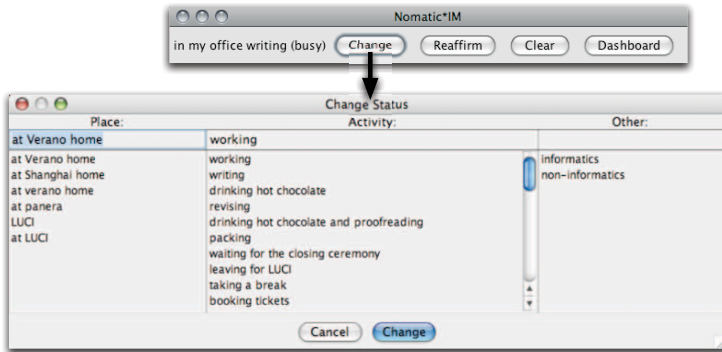


Figure 2. The primary Nomatic window shows the status that is currently being reported to IM systems and Nomatic*Viz. When a user wants to change his status, he is provided with a list of suggestions which are generated by a machine learning algorithm that matches current sensor readings to previously used status messages.

a shared community space. In this section, we will describe Nomatic*IM and Nomatic*Viz in details.

Nomatic*IM uses a wide variety of built-in sensors (e.g. ambient light, motion, network parameters, battery charging status) on a user’s laptop to sense aspects of a user’s context. Unlike IM presence cues, which are almost raw sensor data (e.g., “idle”), Nomatic*IM, whose interface is shown in figure 2, uses machine learning to present a list of predicted status descriptions from which the user may choose. The status messages are combinations of a user’s place, activity and mood and are predicted based on current sensor readings and the user’s history of status entries. Additionally when the system independently thinks that the context has changed or after a period of inactivity (two hours by default), it will remind users to update their status. The goal is to allow users the freedom to richly express their current context without requiring more than a couple of mouse clicks in the best case. The selected status entry is then broadcast to a wide variety of IM systems (e.g., Skype, AIM, Yahoo!, MSN etc.), social networking sites (e.g., Twitter, Facebook) and Nomatic*Viz.

Nomatic*Viz is designed to be situated in a shared community space. Through the design of Nomatic*Viz we sought to leverage ambiguity to address privacy concerns and more importantly to create a thought-provoking and reflective visualization of the entire community’s sensor and status data. In the spirit of work by Gaver et al. (2003) we wanted to focus Nomatic*Viz’s audience on the interpretation of the overall rhythm of the community and not on the specifics of the sensor data available. Our goal was not to engage viewers with the system *per se*, but rather to engage them with the community who is generating the data. By showing status information in an ambiguous way, we hope it will encourage users to relate their contextual social background to more actively interpret the display and experience the community in new ways.

Figure 3 shows a screenshot of the resulting design. All participants’ status data

for the current week is aggregated on the display. On the top, a calendar-like bar indicates to viewers what days' data is currently being shown, the current time of the day is shown with a red arc, and a series of color boxes represents each participant whose data is on the display. Each user has a slightly different color that remains constant over time.

Across the display colorful "fans" are shown. Each cluster of fans corresponds to a unique WiFi access point from which statuses are reported. The arc sweep on the fan corresponds to the time at which the user was at the location (mirroring the clock at the top of the screen). Multiple days of reports from that access point are layered on top of each other. A text label with the most recently reported place name is shown next to the fan. If multiple WiFi access points share the same SSID label (for example many access points administratively managed by the same organization), they are clustered together on a gray ring. The size of the gray ring is determined by the number of unique access points associated with it and its position on the screen is based on the recency of the last status report from it. In the case of figure 3, the central ring represents a university campus. Overlaid across the display, large lines of text representing current activities sampled from all participants appear and disappear. The visualization highlights current real-time status reports with pulsing circles over the relevant fan.

Thus, the information presented in Nomatic*Viz is ambiguous at several levels. Instead of using a literal geographical map as its layout, it is dynamically constructed by users' collective interactions with the WiFi infrastructure. Instead of using icons to represent people, it uses different colors to subtly distinguish individuals. The mapping of colors to individuals is not specified, and is therefore unknown unless viewers have a knowledge of the community. The display does not associate activity text with an individual either. The fact that place descriptions are user generated allow for user control over the degree of accuracy over the location names. Finally, by layering historical traces of people's whereabouts over time details become obscured but frequency becomes more pronounced.

The Study

To understand how people live with Nomatic*Viz in a community setting, we deployed Nomatic*Viz in the authors' academic department, a shared public space, and conducted a five month field study.

The Setting

An entire academic department is located on the floor of the building where the deployment was conducted. To make the display more accessible and to facilitate sharing by the whole community, we placed the display at the entrance lobby to the floor, which was also the connecting point between two wings of offices and is close to many shared resources such as the kitchen, the bathrooms, the copy and mail rooms, etc. Notably, this was not the first spot that we tried. In early




Figure 3. Screenshot of Nomatic*Viz display with four days of data.

iterations of this system the display had been deployed in the elevator waiting area, where, counterintuitively, we found people didn't have time to view the display. The studied display location instead had long sight lines which enabled the display to be viewed while people were in transit to other locations.

The display of the new visualization was mounted in the lobby from the end of March 2008 and remains in place as of the writing of this paper. The preliminary set of 7 participants consisted of researchers affiliated with the Nomatic project (one faculty member, four graduate students, and two undergraduate students). Over the next 10 weeks several other faculty members and graduate students were enrolled through personal invitations. During the summer quarter a new round of participants were recruited to replace participants who were no longer physically located in the building. By the end of the summer, there were 89 users of Nomatic*IM, 16 of which had opted-in to the display of their data on Nomatic*Viz. These 16 participants consisted of 3 faculty members, 1 research scientist, 10 graduate students and 2 undergraduate students.

Methods

After the display was deployed in the field for five months (crossing two quarters and a summer), we conducted semi-structured interviews and analyzed logs of status messages. The interview protocol covered four areas: *everyday schedules and mobile practices*; *interpretation of the visualization*; *encounters with the display* – probing for specific instances; and *disclosure practices*. When possible interviews were recorded and later transcribed. To jog memories and probe specific instances, most of the interviews were conducted by the display, with the interviewer providing a printed samples of past status messages to the informants. Out of the total 16 participants who broadcast status to the display, 8 were interviewed. One interview was with a member of the Nomatic research group, and the remaining 7 interviews were with participants not affiliated with the Nomatic project. Of the 8 participants, 3 were faculty members, 1 was a research scientist and 4 were graduate students. 5 were females and 3 were males. At the time of interviews, all informants had been using the display for at least a month, with several spanning the entire 5 month deployment. In addition to the interviews, logs of status messages were analyzed to gain insight into how they were shared.

Results

Over the period of 171 days, a total of 10772 updates were received from the 8 informants, with an average of 63 updates per day and 4 updates per day per person. Out of all of the updates 471 were unique messages. Figure 4 shows the number of status messages posted per day during the course of the study. A strong cyclical pattern is evident. Figure 5 shows the average number of updates per day-of-the-week. It reveals a strong tendency for people to update status more often early in the week, gradually declining through the week and then much less on




Figure 4. The total number of status updates per day during the duration of the study, with Saturdays identified.

weekends. This trend explains the cyclical update pattern in figure 4 with the low points corresponding to weekends.

Both the length of the study and the relative consistency of status updates throughout the duration of the study suggest that this study was able to mitigate novelty effects. Additionally the consistency of status updates suggest that Nomatic*IM was effective in promoting regular status updates. It was also consistent with findings from our interviews: our informants commented that Nomatic*IM was lightweight, and didn't involve much work to use and therefore made contributing to Nomatic*Viz easy




Figure 5. The average number of status updates per weekday.

as well. Several informants reported that they definitely started updating their status much more. It was especially true for those who didn't update at all through other social media. More specifically, the Nomatic*IM window popping up periodically was effective as a reminder for them to update their statuses.

To provide a framework for what status messages were broadcast by this group, we analyzed status message logs and identified five frequently used themes into which status messages could be loosely categorized:

- **Meeting Events:** Including talks, presentation, meetings. (e.g., “listening to XXX’ talk”, “attending a Ph.D. defense”, “in XXX’s advancement”)
- **Work Activities:** Descriptions of tasks consistent with academic work. (e.g. “hacking”, “coding”, “working on dissertation”, “reviewing papers”, “story-boarding”)

- **Non-work/Leisure activities:** Tasks not associated with the workplace (e.g., “cooking”, “making coffee”, “Watching TV”, “playing games with XXX”)
- **Social Banter:** Information, jokes, invitations, and greetings (e.g., “not in an earthquake zone”, “Happy birthday XXX!”, “early lunch anyone?”)
- **Expressions of Mood:** Emotions, frustrations and reactions (e.g. “WHY DOES EMAIL HATE ME??????”, “punching my computer in the face”, “probably sleeping”, “exploding”)

As part of the evaluation of Nomatic*Viz we conducted semi-structured interviews with the 8 informants. At the highest level, our analysis demonstrated a very sophisticated and nuanced response to the introduction of these status tools into the informants’ daily practice. While a complete understanding of the details of the visual elements of the display was rarely displayed, our informants all developed deep social understandings of the implications of using the tools in various ways. In what follows, we will discuss some of the experiences with respect to the Nomatic*Viz display in particular.

Peripheral Awareness through Glancing

As expected based on prior findings on large displays in shared space (Greenberg and Rounding, 2001; Churchill et al., 2003), the situated context and people’s existing practices shaped how people encountered the display. A great deal of the impact of Nomatic*Viz was related to its location in a lobby that acts a hub connecting the various offices and that is near functional rooms. The lobby itself was not identified as a destination in and of itself and existed on the way to some other pursuit. As a result people were already engaged in an activity when the opportunity for viewing the display presented itself. Frances⁵ noted that most of time when she was passing by, she was busy with other stuff, so she didn’t have the energy to change her actions on the way.



Figure 6. Glancing is the main form of interaction with Nomatic*Viz.

Indeed, in interviews, while our informants reported always “looking” at the display when they pass by, either when they come onto the floor in the morning, or when they visit the copy room, the kitchen, and other’s offices during the day, the majority of them reported that they just “glance” at it. Despite the presence of seating, very rarely would they stop, approach or carefully study it. An exception to this was mentioned by two informants who on occasion would study the display while using the adjacent kitchen to warm their lunch.

Despite just glancing, our data suggests that, although the same data was visible in participants’ IM buddy list status, Nomatic*Viz was still providing a more

⁵With “G” indicating graduate students, “P” pseudonyms whose first letter indicates their role in the department: faculty names start with “F”, graduate students’ and research scientists’ names start

lightweight and peripheral awareness. As Frank put it,

“I don’t really look at my buddy list very often, unless it is just before I am about to make a connection with someone... but if I am walking by the display on my way to the mail room, I just kind of glance at it.”

Fiona reported similar experiences, *“I used it as a lightweight thing to check, as I walk by anyway, while I think about what I’m going to do next.”* This supports previous results by Huang and Mynatt (2003) who point out that the large display, by making information persistent in a shared space, makes information more easily available and relieves the overhead required for people to retrieve it from other channels such as email or IM.

A Community Display

While more lightweight awareness seems to be the most obvious effect of a large display of social status, Nomatic*Viz is distinctive in that it conveys a sense of what is going on within *the community as a whole*, not just as a collection of individuals. This is primarily as a result of the juxtaposition of everyone’s data together, somewhat anonymized and therefore individually obscured.

The most common element noticed while glancing was the rapid appearance of big text. According to Greg, these big words are helpful for *“sampling what is going on”*. George reported similar experiences, *“so when I walk off the elevator, the only thing that I glance at as I walk by are the big things that come up, like the task that people are doing.”* To Fiona, these big words are her favorite feature, *“the number one thing I look at it, which I really like it is the big things that sort of pop by...some sort of ambient knowledge about what is going on with my community of people.”*

Occasionally, the words together formed patterns which characterized particular community contexts such as the end of the quarter, during a conference or a paper deadline, etc. These visible patterns enhanced the feeling of the shared experience in the community. Fiona, being one of the longest members using the system experienced the “ebbs and flow” of community activities through the display,

“Another thing that is sort of amusing to me is the patterns you will see, so towards the end of the quarter, you will see “grading” up a lot, which is often me, Frank and Frances simultaneously, which sort of tickles me, because you get a sense of, OK, the end of the quarter, everyone is grading, or everyone is studying a lot, which is mostly undergraduate students. Similarly, when the [grant] stuff was going on, Frances and I were writing like crazy, so I saw “writing” flash up a lot.”

Frances, another long term user, also noticed the different mix of activities that showed up between the quarter and the summer which was consistent with peoples’ reports that their schedules were more complex during the quarter.

In addition to the big words, the display also more subtly conveys the sense of activity level through its graphical design in features such as “fan” density, the number and distribution of colored dots, the amount of animation, etc. all of which aggregate as status messages are reported during a week. Greg reported how he

perceived the visual cues in conjunction with the big words, “*Just walking by, I [notice] how many blurbs you can see ... high level status messages ... the radar dots that are animated ... how many people are online ... when they’ve been online*”.

However, unless many people were updating on a given day, the sense of community was lost and the display was perceived not as the zeitgeist of a crowd, but simply as an ambiguous display of individuals. Fiona reported a case that was revealing,

“When it was just, for a while, in the early summer, where it is like me, Frank, Frances, that seems about the only people that were there, there were a few people on sometimes, from time to time, .. but most of time...it is sort of interesting, when I go by, I would always know that is going to be one of us... then when there are times, when there are more people, it is nice, it just changes the way that I see it, it just gives me more of the gestalt view of the department, but it is less like personal small group interaction”

Members’ Reflections

In the visual design of Nomatic*Viz, we hoped to create a view of status that provoked reflections on participants’ roles as members of a community. Using ideas described by Sengers et al. (2005) we attempted to leverage user reflection as an important means for us to uncover unconscious values embedded in status broadcasting technologies and related practices. What we found was that the co-existence of multiple audiences (IM and Nomatic*Viz) did encourage reflections, particularly on the effects of the status on different audiences. Fiona described her audience management as follows,

“I used IM for people’s birthday messages, I can imagine that it might be more useful for me to say happy birthday to someone on the display rather than go through my buddy list, and the other way around, there are cases, it makes sense to broadcast to my buddy list, but not to the display.”

Similarly, Greg, who used to post frivolous messages “just to be silly” in social media, was taken aback by seeing his own presence on the Nomatic*Viz display. It made him reflect on the difference between the display and the IM buddy list,

“I posted my status message and then I saw it on the display, I would say, wow that is a personally identifying message, if you know me and what I’ve been up to, you can tell a lot, it makes me consciously aware, what kind of information I disclose.”

Following his reflections, he concluded that “*the display lets the group know what the group is up to, but IM is better for crafting an identity for myself.*”

New Spatial Experiences

Most work on awareness has been based on informational accounts. That is, awareness technologies are presented as a means to inform activities and availability and thus support coordination and communication. Unexpectedly, in our field trial, we

found, many times, people had noteworthy and meaningful moments even when the display was showing information that they already knew. In fact, the display became meaningful when, in a serendipitous moment of shared knowledge, it showed information that they could easily interpret, or when it coincided with their perceptions of the real world space. Our informants described how they would “smirk”, “crack a smile”, “be amused”, “be tickled” and “laugh” seeing some status messages on the display, and thought it was “interesting” and “funny.” A particularly nice example of this was reported by Grace when recalling her experiences with seeing the status message,

“[her message] says, “in a meeting with students” or something like that, I was at the office next door, so I can hear, “oh yeah, she is meeting with students”, so I thought it was interesting.”

In this case, Grace was overhearing her colleague next door having a meeting with students, and at the same time seeing the status message on IM describing the meeting. Although the status message didn’t provide anything new, yet the consistency of it with the real world experience seems to turn a mundane occurrence into something notable and reportable. The meaning of the status, then, is not just in its being descriptive, but rather, together with the space and setting, it participates in engaging sensations, and producing new meaningful spatial experiences.

Perhaps most explicitly, however, this was seen when users saw their own messages. In fact, Frank and Gladys reported one motivation for them to glance at the display was to view their own messages. George also commented that it was funny seeing his own status up there, or others’ status which he can recognize,

“When I see my own status, I kind of smirk. It is just funny, to see. It is kind of fun. I also smirk when I see “making coffee” because I know it was Frank, no one else makes coffee, and I can guess its him.”

Grace expressed similar feelings when encountering the display and seeing her own messages up there,

“Sometimes I laugh because I see my messages up there. Like one day, I had a kind of strange message. I guess it was a couple of days ago, “at work caffeinating”. I put that message up before I went to go get coffee, and here I am at the elevator, it says in big letters “at work caffeinating” while I had a cup of coffee in my hand. I thought it was really funny. I was like, ‘hey, that is mine’ ”

On each of these occasions, the status message explicitly did not provide new information, but, our informants still seemed to be surprised seeing messages of their own. We hypothesize that our participants approached the display as a window into the lives of others. The resulting confusion and internal discord at seeing what was obviously their own near real-time experience was quickly understood and the experience was perceived as humorous. The display introduced a certain gap or suspense between authoring and seeing, ownership and appropriation, and consistency and inconsistency with the real world space. The gap and suspense opened new opportunities to be surprised, and to invoke a meaningful or at least entertaining

experience within space. The value of the display, in this case then, does not lie so much in providing something new, but rather, in its participation and creation of new dramatic and spatial experiences.

Performance

Our informants also considered how they were represented on the display through status message broadcasting. We draw on Goffman's performance framework in forming an understanding of the subsequent sophisticated control of our participants' self-representations.

Goffman uses the metaphor of theatrical performance to examine mundane face-to-face social interactions (Goffman, 1959). In his framework, he used the notion of front, to refer to "that part of the individual's performance which regularly functions in a general and fixed fashion to define the situation for those who observe the performance" (p.22). He further distinguished it into different parts - setting, appearance and manner, and offered the insight that we often expect consistency and coherency among these front parts, and will focus on exceptions to expected consistency among them. In our field trial, we observed similar concerns in maintaining the consistency among different front parts. However, with the presence of digital elements, the front becomes more complex: requiring not just consistency among appearance, manner and setting, but also consistency among digital and physical presence, and consistency among various digital presences.

Maintaining a Coherent Front

As with face-to-face interactions, maintaining a coherent front is part of our informants' considerations when posting their messages through Nomatic*IM. Their desire was to conform to the expectations of the potential audiences' and to avoid unnecessary misunderstanding or wrong impressions. Various strategies were reported by our informants to cope with the complicated and faceted appearances they wanted to maintain. One strategy was to make sure the status message was correct, but not precise. Frank, for example, is explicit about this strategy,

"I don't want people to know specifically where I am, but I let people know I am on campus. If I am running personal errands on campus...I don't want people to think, because I am going down to the store, getting dinner for tonight, or something like that, I am not working on their project."

A second also commonly adopted strategy was for participants to shutdown the whole system, removing their digital presence and to lean on the resulting uncertainty that a complete lack of presence offers to give room for appropriate social interaction. Grace reported an occasion like this,

"There was one week I was sick. I was at home all the time. I kind of didn't want everyone knowing that I am at home all the time, so I didn't put things into Nomatic."

While being at home was not something that was necessarily problematic for this informant, its inconsistency with the usual situations may draw attention and lead to misinterpretations in the absence of further context. By shutting down the system, it saved Grace the trouble of explaining it.

Another common reason for inconsistency is due to a lack of updates or inattention to stale status. For instance, when Frances went on vacation she was more concerned about reporting inconsistent information between the digital and real presence due to inattention than she was about communicating the fact that she was on vacation through the system. As a result, she also shut down the system.

Same Message, Different Audiences

What adds to the complexity of maintaining a coherent front is when multiple audiences are present. It is certainly true with the Nomatic system, where the same status information is published to both the Nomatic*Viz display and IM and others. Fiona is a typical example, as a faculty member, her audience is very mixed, including peers, colleagues, friends, students, people at the distance as well as people that are local. As a response to this very mixed audience, statuses become less funny, and more vague, or innocuous as Fiona put it,

“I make things neutral for the most part. If I want to be funny, I try to make sure it is funny in a way that it is not going to be a problem for the mixed audience that I have.”

To address the multiple audiences, status messages are, to some extent, “washed out”. While managing different audiences separately may relieve some concerns, however, it may not be practical, as Greg points out, “*By dodging the problem, it may introduce new problems.*”

What is particularly striking here, is the sophisticated ways people crafted a single message for different audiences at the same time. One technique was to create *high context* and *low context* (Hall, 1976) messages to communicate different things and suggest different availability to multiple audiences. While the former says things that heavily rely on the shared context to understand what the speaker is saying, the latter incorporates more contextual information in their actual communications. Our data shows a common use of both high context messages and low context messages to say different things to different audiences.

Fiona’s example was a case in point. During the period of the study, she was actively collaborating with a local hospital in another city. Originally, she put “hospital” as her location status. However, this message often invited some questions from her buddies who wanted to check in on her out of concern for her health. It made her change her status message to the city’s name where the hospital is, so that the local people with appropriate context would be informed where and what she was up to, at the same time, it wouldn’t create much concern from remote friends because the city name didn’t make much sense to them without the contextual background. In this case, the city name is a relatively high context message, since it requires the right contextual background to interpret it appropriately.

In reverse, while a high context message is only informative when coupled with shared background knowledge, low context messages are commonly employed to suggest different availability to different people. Grace's example was representative,

"Sometimes, I am working pretty intensively on something. Like, this week, I was working on a NSF grant, and I put my status as "busy with NSF grant", then anyone who is working with me on the NSF grant knows that they can disturb me because they were also working on the same thing. Sometimes that is helpful in terms of gate keeping."

Indeed, our data shows that people respond to these messages accordingly. For example, people will decide whether it is appropriate to interrupt depending on whether the message shown is relevant to their work. In just a very few words Grace was able to rapidly and effectively shift the boundaries of her "groups" in such a way as to invite some contact while preemptively avoiding other types of contact.

We found, sometimes, people customized the status with certain audiences in mind. One extreme example was provided by Fiona. During a stressful time, Normatic*Viz was appropriated to create a little bit of entertainment,

"It must have been around some paper deadlines, that I can't remember, and we were joking around. I was listening to some silly music or something, Justin Timberlake's 'Bringing Sexy Back', so I switched it to my IM message, just to see whether it would show up on the display, to mess around and then Frank and Frances kept running back and forth past the display. They were going between offices. They were often coming in here, and they were writing together, and so I was amused to see whether they would run past the display and happen to notice that that was on there."

By putting a humorous message up Fiona changed the nature of the boundaries of her buddies, rapidly shifting into a mode of inviting contact where previously incidental contact was unwanted.

As an example of a low context message style, Frances mentioned that she liked to craft her messages to communicate, in a way that made sense to the remote audience. For example, when she traveled, she intentionally used the hotel name, instead of the city name as her location status, because to her, that was more indicative of her being in a conference than using the city name.

Discussion

We began our investigation by assuming the messages on the large displays were more public than IM. The framing of large displays in the literature is often that they are the "public" because they present no technical mechanism for restricting viewing. In contrast the laptop is framed as the "private" place where information is kept and carefully managed. In our study however, these categories did not align. Many of our participants had buddy lists that contained hundreds of people. For

our faculty participants the heterogeneous nature of the buddy list was plainly evident. There might be colleagues, family members, former students, administrators, bosses, and children all reading the same status message on IM, Facebook and/or Twitter. Yet in contrast the very fact that the large display was a situated artifact forced it to have a somewhat higher level of privacy. It was extremely unlikely that hundreds of people would see a status message that was put on the display. Many of our participants were comfortable with joking around or even making disparaging comments on the “public” display, but were concerned when they realized that the same message was being broadcast to their “private” IM list. So in what sense then is the large-screen display more public?

Our experiences with Nomatic*Viz and Nomatic*IM in a shared community space suggest that “public” and “private” might not be the best concepts to understand these displays and their use, instead, the idea of different audiences might provide a more useful organizing frame. Our five month field trial revealed that, even with just very simple status editing, our informants put considerable thought into how to present themselves with certain audiences in mind. Our informants all demonstrated sophisticated negotiation processes in terms of revealing more general or more specific status, high context or low context messages, and a digital presence or absence. Each of these choices was based on the sensitivity of the message, a need to invite some interruptions while preventing others, and the need to maintain consistency among different fronts. These observations show that managing status and awareness information is more complex than just concerns about awareness and privacy. It resonates more with the performance framework developed by Goffman who was making sense of everyday social interactions. While Goffman only considered face to face interactions in closed settings, the introduction of presence and awareness technologies into our social milieu creates complications for maintaining coherency and consistency among multiple new digital fronts.

Further, another surprising finding was the observation that our informants found meaning in the display not just in its descriptive nature, but also in the way that it created new spatial experiences. Many times people used words like “laugh”, “smirk”, “crack a joke”, “have fun”, and “interesting” while describing their experiences with their *own* status messages or the status messages whose authors they recognized. These effects could be explained in the way that the display created two distinctive spaces described by Mcgrath in examining surveillance technologies: the space where actions are taking place, and the second space or the “watched space” which does not exist prior to watching and is created by the act of watching itself (Mcgrath, 2003). The existence of the “watched space” introduces aspects of suspense, and a certain open-endedness, which, according to McGrath, creates new ways that space can be experienced and understood. Similarly, the existence of a status display also creates a momentary gap between performance and awareness, and between describing and responding. These moments of suspense and the gaps in the consistency of the digital and physical spaces, allow new ways that space can be experienced. Rather than leaving the awareness messages as self-evident descriptions of events, the display causes the viewers to reformulate their represen-

tation of the circumstances of its authorship and consequentially of the meaning of the spatial environment .

Conclusions

In this paper, we have reported on how people perceive, respond and live with a status broadcasting system, called Nomatic, in an academic community. Nomatic is composed of a context-aware software component, Nomatic*IM, which supports status broadcasting, and a large display, Nomatic*Viz, which shows aggregate status information in a community space. A long term field trial of the system suggests that Nomatic*IM is effective in promoting status updates. It also shows that Nomatic*Viz provides more lightweight and peripheral awareness of community activities than IM or other social software. However, interestingly, whether it represents a “community” or just personal relationships is determined by the number of participants updating in real-time, not just by the design of the display itself. Further, what is particularly striking is how the value of the display does not just lie in simply being informative but also in that it participates and creates new dramatic and spatial experiences. Finally, the study also uncovers how members carefully manage the presentation of their status messages and coordinate this presentation across different broadcast venues: a behavior that resonates well with Goffman’s performance framework.

Awareness technologies have occupied much discussion in CSCW, and as sensors, displays and mobile technologies become increasingly pervasive, we believe more awareness technologies will occupy the space we inhabit, and continually play important roles in enhancing collaboration and connections for distributed as well as collocated groups. Our traditional focus has been on how to automatically sense, reason about and display activity and other contextual information, while at the same time providing enough controls for users to address privacy issues. However, as we can see from our data with Nomatic*Viz, this perspective does not address the complexity of how these displays function in a social environment. Instead, we have seen that people engage with awareness displays with contextual knowledge of the community and the space, and how people actively negotiate self-presentations and maintain coherent fronts to multiple audiences simultaneously. Rather than considering a tradeoff between awareness and privacy, we argue that we should consider the spaces that our awareness displays create and how people perform to different audiences in the presence of the display. Computer-mediated awareness is not simply a given; it is an active process of construction and interpretation, within a social and physical setting and set against a contextual background.

Acknowledgments

Thanks to Gillian R. Hayes for her helpful comments, and to all the anonymous reviews for their constructive feedback. We are also grateful for all the experiences and insights shared by the participants

of our study. This work was supported by the National Science Foundation, award IIS-0713562.

References

- Bly, S. A., S. R. Harrison, and S. Irwin (1993): 'Media spaces: bringing people together in a video, audio, and computing environment'. *Commun. ACM*, vol. 36, no. 1, pp. 28–46.
- Brown, B. A. T., A. S. Taylor, S. Izadi, A. Sellen, J. Kaye, and R. Eardley (2007): 'Locating Family Values: A Field Trial of the Whereabouts Clock'. In: J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang (eds.): *Ubicomp*, Vol. 4717 of *Lecture Notes in Computer Science*. pp. 354–371, Springer.
- Churchill, E. F., L. Nelson, L. Denoue, and A. Girgensohn (2003): 'The Plasma Poster Network: Posting Multimedia Content in Public Places'. In: M. Rauterberg, M. Menozzi, and J. Wesson (eds.): *INTERACT*. IOS Press.
- Cockton, G. and P. Korhonen (eds.) (2003): 'Proceedings of the 2003 Conference on Human Factors in Computing Systems, CHI 2003, Ft. Lauderdale, Florida, USA, April 5-10, 2003'. ACM.
- Dourish, P. and S. Bly (1992): 'Portholes: supporting awareness in a distributed work group'. In: *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA, pp. 541–547, ACM Press.
- Fitzpatrick, G. (2003): *Locales Framework: Understanding and Designing for Wicked Problems*. Norwell, MA, USA: Kluwer Academic Publishers.
- Gaver, W. W., J. Beaver, and S. Benford (2003): 'Ambiguity as a resource for design'. In Cockton and Korhonen (2003), pp. 233–240, ACM Press.
- Goffman, E. (1959): *The Presentation of Self in Everyday Life*. Anchor.
- Greenberg, S. and M. Rounding (2001): 'The notification collage: posting information to public and personal displays'. In: *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA, pp. 514–521, ACM.
- Guzman, E. S. D., M. Yau, A. Gagliano, A. Park, and A. K. Dey (2004): 'Exploring the design and use of peripheral displays of awareness information'. In: *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*. New York, NY, USA, pp. 1247–1250, ACM Press.
- Hall, E. T. (1976): *Beyond culture / Edward T. Hall*. Anchor Press, Garden City, N.Y. :, 1st ed. edition.
- Huang, E. M. and E. D. Mynatt (2003): 'Semi-public displays for small, co-located groups'. In: *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA, pp. 49–56, ACM.
- Hudson, S. E., J. Fogarty, C. G. Atkeson, D. Avrahami, J. Forlizzi, S. B. Kiesler, J. C. Lee, and J. Yang (2003): 'Predicting human interruptibility with sensors: a Wizard of Oz feasibility study'. In Cockton and Korhonen (2003), pp. 257–264, ACM.
- McDonald, D. W., J. F. McCarthy, S. Soroczak, D. H. Nguyen, and A. M. Rashid (2008): 'Proactive displays: Supporting awareness in fluid social environments'. *ACM Trans. Comput.-Hum. Interact.*, vol. 14, no. 4, pp. 1–31.
- Mcgrath, J. (2003): *Loving Big Brother: Performance, Privacy and Surveillance Space*. Routledge.

- Moran, T. P., E. Saund, W. V. Melle, A. U. Gujar, K. P. Fishkin, and B. L. Harrison (1999): 'Design and technology for Collaborage: collaborative collages of information on physical walls'. In: *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*. New York, NY, USA, pp. 197–206, ACM.
- Patil, S. and J. Lai (2005): 'Who gets to know what when: configuring privacy permissions in an awareness application'. In: *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA, pp. 101–110, ACM.
- Patterson, D. J., C. Baker, X. Ding, S. Kaufman, K. Liu, and A. Zaldivar (2008): 'Online Everywhere: Evolving Mobile Instant Messaging Practices'. In: J. McCarthy, J. Scott, and W. Woo (eds.): *UbiComp08*. New York, NY, USA, pp. 64–73, ACM.
- Patterson, D. J., X. Ding, and N. Noack (2006): 'Nomatic: Location By, For, and Of Crowds.'. In: M. Hazas, J. Krumm, and T. Strang (eds.): *Location- and Context-Awareness, Second International Workshop, LoCA 2006, Dublin, Ireland, May 10-11, 2006, Proceedings*, Vol. 3987 of *Lecture Notes in Computer Science*. pp. 186–203, Springer.
- Rowan, J. and E. D. Mynatt (2005): 'Digital Family Portrait Field Trial: Support for Aging in Place.'. In: M. G. Williams and M. W. Altom (eds.): *CHI*. New York, NY, USA, pp. 521–530, ACM Press.
- Sengers, P., K. Boehner, S. David, and J. J. Kaye (2005): 'Reflective design'. In: *CC '05: Proceedings of the 4th decennial conference on Critical computing*. New York, NY, USA, pp. 49–58, ACM.
- Smale, S. and S. Greenberg (2005): 'Broadcasting information via display names in instant messaging'. In: *GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. New York, NY, USA, pp. 89–98, ACM.
- Smith, I. E., S. Consolvo, A. LaMarca, J. Hightower, J. Scott, T. Sohn, J. Hughes, G. Iachello, and G. D. Abowd (2005): 'Social Disclosure of Place: From Location Technology to Communication Practices'. In: H.-W. Gellersen, R. Want, and A. Schmidt (eds.): *Pervasive*, Vol. 3468 of *Lecture Notes in Computer Science*. pp. 134–151, Springer.
- Tang, J. C., N. Yankelovich, J. Begole, M. V. Kleek, F. Li, and J. Bhalodia (2001): 'ConNexus to awarenex: extending awareness to mobile users'. In: *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*. New York, NY, USA, pp. 221–228, ACM Press.
- Terrell, G. B. and D. S. McCrickard (2006): 'Enlightening a co-located community with a semi-public notification system'. In: *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*. New York, NY, USA, pp. 21–24, ACM.

The ‘out-of-avatar experience’: object-focused collaboration in Second Life

Greg Wadley¹ and Nicolas Ducheneaut²

¹The University of Melbourne, Australia ²Palo Alto Research Center, U.S.A.

¹*greg.wadley@unimelb.edu.au* ²*nicolas@parc.com*

Abstract. Much of our current understanding of collaboration around objects in collaborative virtual environments comes from studies conducted with experimental immersive systems. Now Internet-based desktop virtual worlds (VWs) have become a popular form of 3d environment, and have been proposed for a variety of workplace scenarios. One popular VW, Second Life (SL), allows its users to create and manipulate objects. This provides an opportunity to examine the problems and practices of object-focused collaboration in a current system and compare them to prior results. We studied small groups as they assembled objects in SL under varying conditions. In this paper we discuss the problems they encountered and the techniques they used to overcome them. We present measures of camera movement and verbal reference to objects, and discuss the impact of the UI upon these behaviors. We argue that while well-documented old problems remain very much alive, their manifestation in SL suggests new possibilities for supporting collaboration in 3d spaces. In particular, directly representing users’ focus of attention may be more efficient than indirectly representing it via avatar gaze or gestures.

Introduction

Much of our current understanding of collaboration around objects in 3d environments comes from laboratory studies conducted five to ten years ago, with the state of the art defined in books edited by Churchill et al. (2001), Schroeder (2002), and Schroeder and Axelsson (2006). Since then, Internet-based ‘desktop’ virtual worlds (VWs) have become popular environments in the form of massively-multiplayer games and social worlds. These systems use standard PC

hardware rather than specialized I/O devices: limited hardware being both a boon, allowing VWs to achieve mass popularity, and a detriment to the realism of the user experience. There is renewed interest in using online 3d environments in non-recreational contexts such as online meetings and education, with several workplace-oriented VWs in use or under development.

Of the recreational VWs that are currently popular, one of them, Second Life (SL) is exceptional in that it allows users to build and manipulate the contents of the world using 3d editing tools built into the client (Ondrejka, 2005). The emergence of SL as a popular platform provides an opportunity to update our understanding of collaboration around virtual objects by comparing practice in this system with that observed in prior research.

There is current academic interest in VWs at CSCW and elsewhere (eg Brown and Bell, 2004; Moore et al., 2006; Yee et al., 2006), but most of this work has focused on social interaction and communication rather than the mechanics of collaboration in three dimensions. To bridge this gap, we report in this paper on a study of collaborative building in Second Life. We gathered data in several ways. To facilitate comparison with earlier research we conducted a laboratory study inspired by experiments such as that of Hindmarsh et al. (1998). We asked groups of two or three participants to collaborate on building tasks, and recorded their screen video and conversation for later analysis. With each group we discussed the problems they faced, how they solved them, and their thoughts on the user interface. We conducted quantitative measures of participants' deictic verbal references, and their use of SL's detachable camera, an interesting UI feature that potentially impacts collaboration. Finally we discussed themes that arose with other expert users discovered in-world and in the SL forum. We did not undertake traditional online ethnography, because it would be rare to chance upon instances of collaborative building and impossible to see the users' view of the virtual scene portrayed on their screens. However one author was an intensive user of SL during the study and was able to discuss and observe in-world practice in order to ground the laboratory observations.

Collaborative virtual environments

The primary communicative affordance of virtual worlds is a simulated 3-dimensional space in which users are represented to each other as avatars. This allows the simulation of some aspects of offline interaction such as how people position and orient their avatars and how they refer to objects. Research has investigated the mechanics of 'simulated face-to-face' for several years. One thread of inquiry focuses on the similarity and differences between simulated and physical spaces. Yee et al. (2007), for example, demonstrated that Second Life avatars obey real-life proxemic rules. Moore et al. (2006) emphasized that in

current avatar systems, most user actions are not publicly accountable and therefore hinder the micro-coordination of activities.

VWs should be suited to supporting collaborative work in problem domains that are inherently three-dimensional, such as design, repair, and medicine. For example, a technician fixing a machine might converse with remote experts while referring to a 3d representation of the machine, with all parties able to manipulate the model. A significant body of CVE research has focused on collaboration around objects. A central issue has been how one user can deduce another user's point-of-view in order to reference objects deictically: a problem compounded by the limited gestural abilities of current avatars. Hindmarsh et al. (2001) found that groups struggled to achieve common reference to objects even when able to 'point', since users could not always see both the pointing arm and the referent due to the narrow horizontal field of view of desktop systems. Gestures also forced users to spend too much time 'driving the avatar'. Pinho et al. (2002) studied methods of allocating degrees of freedom so that one user moved an object along a ray while another rotated it. Thus collaborators at different positions could combine their viewpoints to place an object efficiently.

Research in shared-video systems has shed light on collaboration, though avatars are not used and the choice of vistas is usually limited to 'scene' or 'head-mounted' cameras. Kraut et al. (2002 and related work) had helper-worker pairs complete a screen-based jigsaw-puzzle. In this arrangement the worker manipulates objects while the helper can offer only verbal assistance. Sharing the scene view, but not the head-mounted views, improved performance, especially when the task was complex and the objects difficult to describe verbally. Goebbels et al. (2003) had pairs collaboratively manipulate a virtual object with the assistance of haptic control and video-conferencing, finding that users spent more time looking at the object than each other except while resolving misunderstandings, and that voice quality was more critical than video.

While many experimental CVEs allowed subjects to communicate by voice, desktop VWs have until recently offered only typed text for linguistic communication. The mechanics of textual turns-at-talk in a VW was studied by Brown and Bell (2004). Text communication during object-focused collaboration was examined in the VW 'ActiveWorlds' by Herring et al. (2003), who found that novices tried to refer deictically to objects, but resorted increasingly to describing them by name. The recent addition of voice-over-IP to systems such as Second Life contributes to making them more 'lifelike' and indeed, recent research shows that it has considerable, though situation-dependent, benefits for coordination of groups in MMORPGs (Williams et al., 2007; Wadley et al., 2007). While VW users engaged in identity-play often prefer to communicate by text, we took it as a given that voice would be used by collaborating workers, and allowed our participants to speak.

Second Life

Since its inception in 2003, Second Life has grown into one of the most popular commercial VWs, with 1.5 million registered users. The SL client uses a standard PC screen for output and keyboard and mouse for input. Since 2007 a voice channel has been included for user communication. SL user accounts may be paid or free: the approximately 80,000 users who pay can own virtual land and build on it. The right to build permanent structures (which are collections of simple shapes called ‘prims’) is the main advantage conferred by a paid account, so this is a reasonable estimate of how many users are building content. A number of authors have commented on SL’s potential as a tool for CSCW: Van Nederveen (2007) proposed it be used for collaborative architectural design, while Rosenman et al. (2006) tested a design system in which SL was supplemented with tools such as a 2D sketch-pad.

The SL client allows users to choose between first-person (through the avatar’s eyes) and third-person (from behind the avatar) views. Unusually, it allows users to move their camera independently of the position and orientation of their avatar, by anchoring it to an object in the local scene. This technique allows a user to gain multiple perspectives more quickly than is possible by walking an avatar around an object: thus it is commonly used for object-related activity such as building and looking at other users’ creations. Unlike most VWs which support limited camera movement near the avatar, the SL camera can be moved over a wide area, oriented in any direction including up and down, zoomed a long way in and out, and unlike avatars, moved through objects. While users’ avatars are publicly visible, their camera positions are not (Irani et al., 2008). The ability to decouple one’s camera from one’s avatar is similar to techniques suggested by Hindmarsh et al. (2001) and Bailensen et al. (2006).

At any given moment an SL user is viewing the virtual landscape from either their avatar location or the location to which they have moved their camera: thus an SL user’s presence is divided between two different locations. We call these modes ‘in-avatar’ and ‘in-camera’ to emphasize that while avatar locations are visible, camera locations are private. These are illustrated in figure 1. Since the objects that a user can interact with are those that are currently visible to them, rather than those that are in the vicinity of their avatar, it is their private camera-location rather than their public avatar-location that defines their focus of attention.

The detachable camera feature means that there is no reliable relationship between what an SL user can see and what their avatar appears to be looking at. While the feature is useful while editing, prior research (e.g. Hindmarsh et al., 1998) suggests that when users cannot deduce each others’ vistas, their ability to collaborate is lessened.

Second Life provides no specific support for collaborative building beyond the ability to visualize one's collaborators in the shared workspace and to communicate using text or voice. Users can display a map of the local area showing avatar locations and the outlines of buildings; however the map is not sufficiently detailed to assist with object manipulation and tends to be used only for coarse navigation. There is no analogy to peripheral vision in the SL display. Limited pointing is possible: when a user is editing an object, their avatar's arm reaches toward the object and a dotted line (the 'selection bar') connects arm and object. This is similar to the line provided by Hindmarsh et al., the rationale for which was that moving an object at a distance represents projection beyond the avatar. In SL this line provides a rough indication of which object a user is editing; however if object and avatar are sufficiently far apart it is difficult for others to follow the selection bar between them. The editing user sees a highlight on the object, but this is not visible to others.

Given the current popularity of virtual worlds and interest in them as platforms for CSCW, we perceived an opportunity to update the understanding of collaboration in 3d by observing it in a current VW. We chose Second Life because it has a significant user base, a focus on social interaction, and offers object manipulation via in-built tools.

Methods

We used methods based on Hindmarsh et al. (2001) and Kraut et al. (2002) to conduct a 'quasi-experiment' in the sense of Hindmarsh. Participants logged into Second Life in groups of two or three and collaborated on building tasks. Groups were co-present in our lab and arranged so they could hear but not see each other. We observed the groups and recorded their screen output and voice conversation for later analysis. Building sessions were followed by focus-group discussions.




Figure 1: Two users' vistas, captured simultaneously, just before completing the House task. The user on the left is 'in-avatar'. The user on the right is 'in-camera' and editing the roof.

Each complete session was designed to last an hour and a half, although some enthusiastic groups worked over time, which we allowed. We observed ten groups; a total of 22 participants. Their ages ranged from 20 to 50. Half were male. Some knew each other before undertaking the task, while other groups were meeting for the first time. We chose participants with a broad range of SL experience, reasoning that while expert users would demonstrate cutting-edge practice, it was important also to understand novice use. Two participants' recordings were excluded due to technical faults: for statistical tests this left a cohort of four experienced builders (recruited via Craigslist and the SL forum), ten users with experience of other VWs or modeling tools, and six novices.

After spending a few minutes familiarizing themselves with each other and the lab setup, group members undertook two tasks. The first adapted the 'furniture world' task of Hindmarsh and the 'jigsaw puzzle' task of Kraut (though in 3d). The researchers provided a set of virtual objects including four walls, two roof parts, two gables, a flag, flagpole and chimney. The group had to assemble these into a house. Some of the objects, such as the flagpole, were unique, meaning they could be referred to by name. Others, such as the walls, were indistinguishable from each other, requiring spatial reference to identify them. Each group member received a screen-shot of how the house should look when complete. The 'House task' was unstructured, allowing groups to discover their preferred method of collaboration, naturalistic in that dwellings are a popular building project in SL, and relevant to remote repair in being a complex object made of smaller parts.

In pilot studies we found that some users chose to collaborate with minimal communication (more on this later), and so we provided a second task, the 'Garden task', using the helper-worker arrangement of Kraut and designed to force closer collaboration. One group member was designated the helper, while the other(s) were workers. Helpers could not use the building tools themselves. They received a screen shot showing the house now surrounded by extra objects such as garden furniture and a fence, and had to direct their worker(s) to build this scene.

After the two tasks were complete we conducted a focus-group in which we asked participants about their experience, the problems they encountered, how they solved them, and whether the UI could be enhanced to better support collaboration. Focus-groups were semi-structured to allow exploration of themes.

We used the screen recordings to conduct two quantitative analyses of group performances, counting 'salient events' (Schroeder et al., 2006). We measured the proportion of time users spent with their camera decoupled from their avatar, and we categorized verbal references to objects by their linguistic form. We compared these measures across experience levels and tasks to discover patterns of use.

As a final check we discussed emergent themes with a range of experienced users discovered on the SL discussion forum and at building classes and competitions.

Results

Our participants displayed a broad range of building and communication styles, illustrating the variation that general-purpose systems need to support. They encountered communication problems of the kind discovered in earlier research, but worked around them by experimenting with different referential practices (as in Hindmarsh et al., 1998) and rarely became mired in problems of reference.

The following exchange illustrates the kinds of problems that arose and the variety of reference techniques attempted. This group of three is deciding how to position four walls to form the base of their house. The walls look identical and are currently positioned randomly around the vicinity.

A. So are you putting the walls together?

B. I'm moving one wall ... a third wall, towards the other two ... the one that's tilting. [*B marks a wall by changing its orientation in a fashion visible to her team-mates.*]

C. Oh, that was you! [*A and C now know which wall B is editing.*]

A. Why don't you turn that over, and I'll move the other wall? [*'That' refers to the tilting wall, and 'other' to the fourth wall not discussed yet.*]

C. Are you moving the one on the lower leftmost of the walls?

A. Well .. your left? [*A and B laugh, because C attempted a spatial reference which A and B cannot decode.*]

A. I'm going to move the one that I'm standing right next to. [*A doesn't attempt to correct C's attempt at deixis, but instead moves her avatar beside a wall to mark it to the others.*]

A. The one that's I guess kind of closest to [C] ... [*A uses the position of C's avatar as a reference point*] ... why don't we leave that one still, and then we can put the other three around it?

C. Mine's above the ground. [*C refers to the wall closest to his avatar as 'his' wall.*]

A. That's fine I think. Why don't we just leave that one and put the other three around it?

The group tried different referential methods until one worked and they could proceed with the task. Participants often used their avatars to mark positions, such as this exchange from a helper-worker pair performing the Garden task:

H: You see this table I'm standing next to? Don't move this one – this one stays in place.

W: Yep. So why don't you just move where you want the others.

H: Yeah. [*walks to a different spot*] The other one is going to go here. In front of me.

W: Right, hang on ... [*W moves the table*]

H: And the last one is on the other side ... [*walks around the house*]... Just about here.

Another pair's exchange illustrated several techniques: pointing with the edit bar, marking places with avatars, and verbal description of an object:

X: Is there a way to point? What's the thing you thought was the flagpole?

Y: Hang on, let me just walk into it. [*walks to the flagpole*] See this thing that's right near my hand? [*He is currently editing a wall – a different object - so his arm is in the air.*]

X: Which hand?

Y: Right in front of me. Can I point at it? [*He places the flagpole in edit mode so as to point directly to it.*] There we go. Why don't I move it. If you're watching it, I'm moving it back and forwards now. Can you see an object that keeps moving left and right?

X: Yes that's the flagpole isn't it? [...] You just walked past a cement block. Are we supposed to do something with that?

Y: I think that's the chimney.

Experience made a clear difference to participants' ability to collaborate. Yet all groups were able to complete their tasks, albeit at different speeds and with varying quality.

Use of the detachable camera

The ability to rapidly gain multiple perspectives of an object by detaching one's camera from one's avatar is not available in most VWs. While this feature supports efficient building by individuals, it also breaks the relationship between an avatar's orientation and what the avatar's owner can see. Since prior research indicated that deducing collaborators' viewpoints was a significant problem in CVEs, we were interested in how often users detached their cameras while working in SL and whether this affected either their ability to communicate or their experience of virtual embodiment.

Referential problems caused by the mismatch between avatar and camera viewpoints are illustrated by the following exchange. This pair performed their tasks well, but were plagued by an on-going misunderstanding over viewpoints, because one (here called 'A') stayed mostly in-avatar while the other ('C') had his camera zoomed out and seemed to ignore the avatars. Here they have assembled four walls and are about to place the gables. Their avatars stand at opposite ends of the house, however C's camera is near A's avatar, so that unbeknownst to A they are viewing from the same side of the house.

C: Let's place those triangle things. [*the two gables*]

A: Where are those? Oh, the triangle things are around the front aren't they? [*It is not clear which end of the house is the front.*] I'll place the one on my side if you place the one on the other side.

A: I don't know whether I've selected the same one as you. I'm selecting the one that's further from the house.

C: Ok, do you see one moving? I selected one that I just raised up.

A: Yes I see that one, ok good. I'll pick a different one then. Oh you're putting it on that edge?

C: I put it on the nearest spot I could find. *[They both intended the end nearest their views.]*

A: Where's your character? *[he means 'avatar']* Oh ok, I see where your character is. I tell you what, can you put the gable on the house section closest to you, and I'll move the one that's closest to me? Unless you want to finish placing the one that you had. *[A is still using avatar-relative reference while C assumes he means relative to camera.]*

C: Does it matter? I'm maneuvering the one that I had.

A: Ok, I can move the other one I think. I'll just walk around so I can see it better. *[He walks his avatar to the other end of the house, where C's avatar, though not his camera, is placed. While walking he apologizes for bumping into C's avatar, though C was unaware of it.]*

At one point during the Garden task this pair tried the house's frame of reference, and then spatial deixis, before being successful with avatar marking:

C: If you're facing the front of the house, you need one table in the front of the house with two chairs, one to the left of the house with four chairs, and one behind the house with two chairs.

A: So we're going to treat me as facing the house right now? Do you want to see where I am?

C: Um, I see where you're facing.

A: I tell you what, can you walk your avatar to what you're calling the front of the house?

[C goes back into avatar and walks to the front]

A: Ok. So you're currently at the front of the house?

C: Yes I'm facing the front. *[A proceeds to place the furniture.]*

Some experienced MMORPG users rarely detached camera from avatar, suggesting that extensive gaming experience may make disembodied viewing in 3d systems feel unnatural. Experienced SL users said that maintaining both an avatar and a camera location did not bother them. On being questioned about "being in two places at once", most said that this had never occurred to them. When asked, "How would you describe your location right now?", experts usually chose their avatar rather than camera location. One felt that the detached camera was simply a tool, and that while using it he continued to equate his avatar with himself. Conversely, others felt that, while they were building, their avatar was irrelevant and even got in the way. One expert said that as a beginner she had identified with her avatar, but that over time she had begun to experience SL more as a building tool than a virtual reality. But her equally experienced building partner felt that his avatar mattered because avatar locations are how people find each other, and because: "something can happen to your avatar. You can get pushed or shot. Nothing can happen to the camera: it's just a view of a picture."

We calculated how much time each participant spent 'in-avatar' and 'in-camera'. Averaged over all participants, about half of task time was spent in each mode. Suspecting that this correlated with experience, we classified participants and compared their camera use. Group A (n=4) had significant SL experience, while group B (n=10) were competent users and group C (n=6) were novices.

ANOVA showed expertise to have a significant effect on camera use ($F(2,33)=8.93$, $p<.001$), with expert SL users, as expected, more inclined to detach their camera (figure 2).

Task (House vs. Garden) was not a significant factor in camera use, and there was no interaction between expertise and task, however this may represent two effects canceling each other out. In the Garden task, helpers often used their avatars to mark locations, and were in-camera less often. But some workers spent more time in-camera during this task, perhaps due to increased familiarity with the UI.




Figure 2: Time spent 'in avatar'.

Verbal reference to objects and places

Natural language offers several ways to refer to objects and places. CVEs afford spatial deixis because users are embodied at particular locations and orientations which are visible to other users. If a user is looking through the eyes of his avatar, then a deictic reference relative to the avatar is also relative to the user's vista, and should be understandable by him. Therefore a detachable camera should make deixis less reliable.

We examined our participants' use of verbal reference, counting references to objects and locations and categorizing them according to the frame of reference used. A number of categorizations are available: we used that of Levinson (1996), who recognizes relative, intrinsic, and absolute frames. A 'relative' reference involves deixis from the speaker's or receiver's point of view and is the key form of interest here. Some objects have their own 'intrinsic' frames of reference, for example houses may have an obvious front and rear to which other locations can be compared. Finally locations can be relative to an 'absolute' frame of reference such as compass points or a prominent object in the distance. We added a fourth category, 'reference by name or property' to count references such as "the brown

rectangle”. Other researchers have used slightly different schemes: for example Herring et al. (2003) categorized references as “deictic”, “fixed unique” and “fixed non-unique”. Their ‘deictic’ category corresponds to our ‘relative’ category, while their ‘fixed unique’ corresponds to our ‘name or property’ category. We did not count how many of the instances of deixis were successfully interpreted, as this is not always clear to an observer.

Although both Cartesian (x y z) and cardinal (north south east west) frames are available in Second Life, they were rarely used in our study. On only two occasions participants made use of Cartesian coordinates to describe locations. Only one group used the cardinal frame. Although landmarks were visible, on only one occasion did a participant use one for spatial reference, describing a wall as “the side closest to the sea”. The ‘absolute’ category is excluded from the graph below.

Figure 3 illustrates the relative frequency of these forms of reference. These were consistent across groups ($F(2, 27) = 6.37, p < 0.01$). Neither task nor expertise level were significant factors, though expertise affected the overall number of references, with experienced participants making more. While this appears to contradict a finding of Kraut et al. (2002), it suggests that those who were better able to handle SL’s particular style of representation were more comfortable communicating about location. Some novices seemed to be so focused on grappling with the building tools that they neglected to communicate with teammates.




Figure 3: Frames of reference used to refer to objects in speech

It is noteworthy that even experts used deixis relative to a collaborator’s avatar despite knowing the collaborator was probably ‘in-camera’. This was not a problem for experienced users, who seemed able to use their avatar’s frame of reference even while their visual focus was elsewhere. Experts interpreted a reference such as “to your left” in the only reasonable way, meaning “to the left of your avatar”. (Sometimes reference was explicitly relative to “where your avatar

is standing”).) If their avatar was out of view, they moved their camera to bring it into view, pressed Escape to return their camera to their avatar, or asked for more explanation. Expert informants confirmed that they were able to interpret avatar-relative deixis, the only problem being that, if they had to return to their avatar, they lost their camera position and would have to find it again. In Second Life, avatars are permanent cursors that mark a location that is easy to return to, while camera locations are invisible and impermanent.

Occasionally, as illustrated in the exchanges above, novice participants forgot that their avatar but not their camera position was visible. But in most cases, deducing collaborators' vistas did not seem to be a major concern. No participant ever asked a team-mate, before using deixis, whether they had detached their camera. We asked participants whether they would like an extra screen displaying their collaborator's vista, but none thought this would be useful, though some novice users proposed a 'see what I see' feature that could transmit their vista to a colleague when required.

Non-verbal reference

Referencing an object by pointing to it (selecting it for editing) was only rarely observed in our study, though one expert claimed that this technique was used "all the time" in SL. Some participants remarked that although they could see a colleague's avatar point while they were editing, the highlight on the object was only visible to the editor, which reduced the usefulness of pointing as a means of collaborative reference. Avatar gestures other than pointing were never used by participants. Moving one's avatar to stand beside an object however was frequently used. Some participants jiggled objects back and forth or moved them to mark them for collaborators, as illustrated above and in the following exchange:

A: Are you rotating one of your walls?

B: Yes, is it rotating on your screen?

A: I'm rotating another wall. Yep, I just saw yours rotate.

One expert changed the colour of a wall to mark it: however she said this was not a common technique because it is usually hard to restore the original texture.

Some experts said that it was common practice in SL to create a prim to mark a location, to be deleted later after use. However no-one did this in our trials. It is possible that the technique did not occur to novices, and that experts did not find our tasks sufficiently difficult to require it.

Problems editing in 3d

Although not all participants were regular users of 3d environments, none of them had difficulty navigating their avatar around the space. However those unused to SL found it difficult to manipulate objects, which are constrained to move along orthogonal axes. Novices found translating and rotating objects frustrating, some stating they would prefer to simply drag an object from one location to another rather than execute separate moves along each axis. One participant typed destination coordinates rather than use the mouse. An environment for supporting remote repair might require more natural movement of objects.

Prior experience with modeling packages substituted for experience with SL's building tools. One such participant felt that SL was "Blender grafted onto a game". Experienced users sometimes performed an entire build 'in-camera' and on completion walked around it 'in-avatar', suggesting that they experienced building as a distinct activity within the VW.

Novices seemed to assume that objects were subject to collision detection and gravity, though they clearly were not. Experts made use of the fact that objects could hang in the air and pass through other objects, and had a natural orientation along the world's axes. Experts did not express interest in improving the building UI, though several complained that the permissions system was not conducive to collaboration.

Group organization and division of labor

Groups performing the House task were free to organize themselves any way they wished. Different groups divided their labor at different levels of abstraction. A striking proportion chose one method, which was to decompose the house into 'base' and 'roof' sub-assemblies, to be completed separately by individuals and joined in the final stage. This appeared to be a strategy for reducing the need for close collaboration. Most research on collaboration around virtual objects has focused on closely-coupled collaboration, yet we rarely observed participants choosing to work simultaneously on the same prim.

Expert informants reported that although team building is common in SL, especially on large, complex projects, closely-coupled collaboration at the level of individual prims is rare. Experts find it more efficient to decompose a project into sub-tasks, allowing specialization and schedule independence. For example, one user might create a building's skin, while another creates its furniture, another builds walkways and a fourth applies textures. Often these components are not even built at the same site, but are created on the individuals' own land and moved into place in the final phase of building.

Apart from reducing the need for coordination, another explanation for SL users' disinclination to collaborate closely might be that is only necessary in systems which tie vistas to avatars. In those, a user cannot easily obtain multiple

viewpoints of an object and might benefit from feedback from colleagues located at different viewing angles, in the fashion described by Roberts et al. (2006). SL's camera reduces the need for this, and we observed only two occurrences.

Close collaboration in SL is also made less effective by the way visual feedback on object movement is shared. While a user is dragging an object they receive visual feedback at all positions along the object's path. However collaborators only see the end point of the movement. This is probably designed to reduce the number of scene-update messages sent over the network, however it makes it harder for one user to guide another user's placement of an object.

Discussion

Comparing old and new

Our purpose was to observe how the problems and practices of collaboration around virtual objects have evolved since CVEs emerged from research labs to become mainstream technologies. We chose Second Life as our study system because it is a popular VW that allows users to manipulate objects. We conducted a semi-naturalistic exploratory study rather than a formal experiment, but exploited the lab setting to conduct two quantitative analyses.

Many of our findings can be directly compared to previous work. Hindmarsh et al. experimented with extended pointing, peripheral vision and a plan view. SL provides such an extended pointing via the 'selection beam', but we heard mixed views on its usefulness. Instead, some users said it would be more useful if the object highlight seen by an editor was also visible to others, so that knowing who is working on what was more transparent. In other words, there seems to be little need to tie selection feedback to the avatar: changing the visual appearance of the selected object is enough. This seems to contradict (for object-focused work) Moore et al.'s argument that avatar systems should be made richer to enable tighter coordination. In fact it could be argued that representing user's bodies is superfluous during this form of collaboration. Certainly some participants forgot about their avatars, or used them as object-marking cursors.

Related to this, Hindmarsh et al. reported that stylized gestures were not useful for collaboration around objects, and this is supported by our study, which observed no use of SL's pre-programmed gestures. With regard to peripheral vision, it is not provided in SL, and most participants felt that extra screens would be a burden. SL offers a plan view (the map) but none of our participants used it, probably because it displays insufficient detail about objects to be useful.

Hindmarsh et al. found that avatars often gave the wrong impression of what their users could see. SL's detachable camera would seem to make this problem intractable. Yet our participants were able to communicate successfully, if at

times slowly, about objects and locations. Moreover they stated that an extra window showing their collaborator's camera view would be superfluous. This accords with the finding by Fussell et al. (2003) that collaborators preferred seeing the shared workspace to looking through the head-mounted camera of a collaborator (see also Hindmarsh et al., 2001: p134-5). Herring et al. (2003) found that novice users of ActiveWorlds ceased attempting spatial deixis when they found their colleagues could not dereference it. Our participants were more successful with deixis, which may be due to better graphical representation in SL, or the availability of voice, which is better for quickly resolving ambiguity (Löber et al., 2006). Alternatively it may reflect an increased familiarity over time with the representational style of VWs.

The 'out-of-avatar experience'

Second Life is unusual among avatar-based systems in allowing users to detach their camera from their avatar. This feature trades the benefit of rapid acquisition of multiple viewpoints against the drawback of making some deictic references irresolvable and the foci of activity invisible. SL users effectively have two locations, their (public) avatar position and their (private) camera position, a situation which can foster deception (cf. Irani et al. 2008) and conceivably, a reduced sense of embodiment in the avatar.

Experts seem to maintain a sense of where their avatar is facing, and use this to dereference their collaborators' deictic references. When this is impossible they simply return their camera to their avatar to translate the reference. This costs only the time taken to subsequently return to the camera position, and does not seem to significantly impact users, except for a temporary loss of 'state' – they have to remember where their camera was and manually re-establish their view. This difficulty could be easily addressed by adding a 'toggle' to switch back and forth between the two perspectives. Supporting such an ability to smoothly transition between various states and viewpoints might be a fruitful avenue to explore by future VW designers, especially if they intend to support collaborative activities.

It is possible, using scripting, to provide the location and gaze direction of one's camera to other users. This can enable 'you see what I see'. This feature is not available in SL's standard UI, but has been implemented by one entrepreneur as a software add-on which is available for purchase. We obtained this but found it had limited utility. Only one user could send their camera position, and their collaborators could only receive. Switching between one's own camera and the sender's was slow and awkward. Only the camera position was transmitted, not other screen visuals such as editing highlights, thus masking much of what the sender was doing (cf Irani et al., 2008; Moore et al., 2006). One of our participants, an experienced SL builder, was familiar with this add-on but did not use it in his work.

One might expect a detachable camera to diminish the relevance of avatar location. Despite this, researchers have found that SL avatars obey physical-world proxemic norms of inter-personal distance and eye-gaze (Yee et al., 2007). For proxemics to work, users must perceive each other as having a definite location and orientation. In our study, participants often ‘parked’ their avatar while building, moving it to a socially appropriate position only when interacting with other users. In one session, a participant stayed in-camera except when a new user appeared nearby, whereupon he went in-avatar and walked over to them. It seems that SL users regard their avatar as a mediator of social interaction which can be ignored while editing objects. However in collaborative building, which is simultaneously object-focused and social, these two attitudes contradict.

People communicating in SL often place their avatars face-to-face, even if they also detach their camera. Arguably it is a form of ‘perception management’ to maintain proxemic norms with one’s public embodiment while one’s private focus is elsewhere. A user who maintains a conversational orientation while moving their visual focus must be aware that other users may be doing the same thing. A solution to this ambiguity might be to display camera positions on screen. An option in the SL client’s ‘Advanced’ menu allows one to see the locations of nearby cameras; however these are not labeled with avatar names - thus users can know that they are being examined, but not by whom. Cameras are often moved so quickly that to keep track of them is cognitively difficult and would require increased network traffic.

It is sometimes argued that generations who have grown up with 3d videogames will readily adapt to collaboration in virtual environments. But participants with game but not building experience reported difficulties using the building UI, which was perceived as being unaligned with the avatar UI.

Articulating collaboration

Using the terminology of Schmidt and Simone (1996) we can analyze SL as a CSCW system for creating the virtual world’s contents. During collaborative building the common field of work is objects and the virtual space within which they reside. SL provides no specific mechanism for articulating work beyond its regular communication tools.

We were surprised to find that the style of collaboration we have referred to as ‘closely-coupled’, in which two users work on the same primitive object at the same time, was rarely performed. On the contrary, the first impulse of many groups was to modularize their task. Users seem to have devised organizational processes that preclude the need for fine-grained collaboration, and there may be several reasons for this. One is that the articulation work required for close collaboration in a 3d environment represents too high a load. Another is that it might be easier to gain multiple viewpoints by moving one’s camera than

receiving verbal feedback from a collaborator. A third is that SL's permission system forces one to explicitly change a default setting in order to allow collaborators to edit objects one has created. A division of labor that involves individual construction of separate modules seems to better leverage the benefits of having more than one person involved in the task.

We implemented a helper-worker task in order to encourage more communication about objects and location. It is noteworthy that in other research where participants worked closely around individual objects, close collaboration was also 'forced'. For example, Pinho et al. (2002) required one user to move an object which was distal to their avatar, while another user closer to the object guided its placement. Roberts et al. (2006) implemented gravity so that two users were required to lift objects while a third joined them together. By contrast SL allows users to rapidly acquire a variety of viewpoints and does not implement gravity by default, so that objects can be lifted by a single user and will stay in place while the user works on other objects.

It may be that VW users will only collaborate closely around objects if physical-world constraints such as gravity and strict embodiment of camera within avatar are reintroduced. But these constraints are not necessary in a virtual environment. Rather than insisting on mimicking physical reality to encourage tightly-coupled interaction, it seems more productive to embrace a VWs' 'unrealistic' properties. As an example of this dichotomy, we would cite again the suggestion by Hindmarsh et al. that users should be made aware of their collaborator's viewpoint (mimicking the accountability of actions from the physical world, cf. Moore et al., 2006) and compare it to a possibility suggested by our study, namely, that users should be able to switch at will between several viewpoints. The 'unrealistic' ubiquity we propose might turn out to be more productive than insisting on reproducing the more familiar, but ultimately more limiting, 'one body – one view' paradigm.

Conclusion

The appearance of Second Life, a popular Internet-based virtual world that allows users to edit its contents, provides an opportunity to update our understanding of collaboration around virtual objects. It is interesting to note that problems identified more than ten years ago in experimental CVEs are still prevalent in a 'mass market' environment like SL. In particular, difficulties with the UI (especially the lack of transparency and feedback about a collaborator's actions) can lead to a tendency to partition collaborative building into isolated, individual sub-tasks that can be completed in parallel and assembled only at the very end. But our users did not react positively to suggestions from past research that could have made tightly coupled collaboration easier. Shared viewpoints, for instance,

were considered to be cumbersome and unnecessary; avatar gestures for pointing were rarely used; etc.

Instead, our data suggest another avenue for supporting collaboration in VWs: ‘decoupling’ them from physical reality to leverage their unique properties. For instance, while users frequently did not use their avatar like we would use our bodies in physical-world collaboration (by pointing, orienting to the object, etc.), they asked instead for the object itself to be more accountable: for instance, making visible the fact that it is selected by someone else, relaying movement as it happens rather than only at the end of a sequence of modifications, etc. It is technically easy to make the state of an object visible in VWs and yet they have remained for now silent partners in collaborative tasks. We argue that a lot could be gained by thinking about how to make objects, rather than avatars, richer and more interactive.

In a related fashion, the separation between ‘in-avatar’ and ‘in-camera’ modes did not introduce as many coordination problems as one might have expected – and in fact, experts used the two modes to conveniently handle relative positioning and object manipulation synchronously. Rather than trying to reconstruct a collaborator’s ever-changing viewpoint, it might be more productive to accept that users can literally be in several places at once and instead make the transition between various modes more straightforward. As we saw, users will then switch to whatever viewpoint is necessary for the task at hand without much misunderstanding of what their collaborator says.

The fact that, while building, many users ‘parked’ their avatar and concentrated on the objects instead suggests that while it is necessary to represent users’ focus of attention for collaboration around objects, it is less necessary to represent their bodies. Indicating attention directly via a shared cursor or by highlighting objects could be more efficient than indirectly representing it via an avatar’s eye-gaze or gestures.

As collaborative VWs become more mainstream it will be interesting to see whether the practices we observed are truly widespread. In the meantime, we hope this study will inspire VW designers to explore more ‘unrealistic’ interfaces to better support collaboration in 3d spaces.

Acknowledgements

We would like to thank our study participants, and especially Don Wen, Diane Schiano and Mike Roberts at PARC, Jonas Karlsson at Xerox, and Martin Gibbs at the University of Melbourne for their help with this study.

References

- Bailenson, J., Yee, N. and Merget, D. (2006) 'The Effect of Behavioral Realism and Form Realism of Real-Time Avatar Faces on Verbal Disclosure, Nonverbal Disclosure, Emotion Recognition, and Copresence in Dyadic Interaction', *Presence: Teleoperators and Virtual Environments*, vol. 15, no. 4, August 2006, pp. 359-372.
- Brown, B. and Bell, M. (2004) 'CSCW at play: 'There' as a collaborative virtual environment', in *Proceedings of the 2004 ACM conference on Computer supported cooperative work*. Chicago, pp. 350-359.
- Boellstorff, T. (2008) *Coming of Age in Second Life*, Princeton University Press, Princeton.
- Churchill, E. F., Snowdon, D. N. and Munro, A. J. (eds.). (2001) *Collaborative virtual environments: digital places and spaces for interaction*, Springer, London.
- Fussell, S. R., Setlock, L. D. and Kraut, R. E. (2003) 'Effects of Head-Mounted and Scene-Oriented Video Systems on Remote Collaboration on Physical Tasks', in *Proceedings of the SIGCHI conference on Human factors in computing systems*, Ft Lauderdale, pp. 513-520.
- Goebels, G., Lalioti, V. and Göbel, M. (2003) 'Design and evaluation of team work in distributed collaborative virtual environments', in *Proceedings of the ACM symposium on Virtual reality software and technology*, Osaka, pp. 231-238.
- Herring, S. C., Borner, K. and Swan, M. B. (2003), 'When rich media are opaque: Spatial reference in a 3-D virtual world.' Invited talk, Microsoft Research, Redmond.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S. and Greenhalgh, C. (1998) 'Fragmented Interaction: Establishing Mutual Orientation in Virtual Environments', In *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, New York, pp. 217-226.
- Hindmarsh, J., Fraser, M., Heath, C. and Benford, S. (2001) 'Virtually missing the point: configuring CVEs for object-focused interaction', in Churchill, E. F., Snowdon, D. N. and Munro, A. J. (eds.): *Collaborative Virtual Environments: Digital places and spaces for interaction*, Springer, London, pp. 115-139.
- Irani, L. C., Hayes, G. R. and Dourish, P. (2008) 'Situated practices of looking: visual practice in an online world', in *Proceedings of the ACM 2008 conference on Computer Supported Cooperative Work*, San Diego, pp. 187-196.
- Kraut, R. K., Gergle, D. and Fussell, S. R. (2002) 'The Use of Visual Information in Shared Visual Spaces: Informing the Development of Virtual Co-Presence', in *Proceedings of the 2002 ACM conference on Computer supported cooperative work (CSCW 02)*, New Orleans, pp. 31-40.
- Levinson, S. C. (1996) 'Frames of reference and Molyneux's question: cross-linguistic evidence', in Bloom, P., Peterson, M.A., Nadel, L. and Garrett, M.F. (eds.) *Language and Space*, MIT Press, Cambridge.
- Löber, A., Grimm, S. and Schwabe, G. (2006) 'Audio vs chat: Can media speed explain the differences in productivity?', in *Proceedings of the 14th European Conference on Information Systems*, Goteborg, pp. 2172-2183.
- Moore, R., Ducheneaut, N. and Nickell, E. (2006) 'Doing Virtually Nothing: Awareness and Accountability in Massively Multiplayer Online Worlds', *Computer Supported Cooperative Work*, vol. 16, no. 3 pp. 265-305.
- Ondrejka, C. R. (2005) 'Escaping the Gilded Cage: User Created Content and Building the Metaverse', *New York Law School Law Review*, vol. 49, no. 1, pp. 81-101.

- Pinho, M. S., Bowman, D. A. and Freitas, C. M. D. S. (2002) 'Cooperative Object Manipulation in Immersive Virtual Environments: Framework and Techniques', in *Proceedings of the ACM symposium on Virtual reality software and technology (VRST 02)*, Hong Kong, pp. 171-178.
- Roberts, D., Wolff, R. and Otto, O. (2006) 'The Impact of Display System and Embodiment on Closely Coupled Collaboration Between Remote Users', In Schroeder, R. and Axelsson, A.-S. (eds.) *Avatars at Work and Play: Collaboration and Interaction in Shared Virtual Environments*, Springer: London.
- Rosenman, M., Merrick, K., Maher, M. and Marchant, D. (2006) 'Designworld: A Multidisciplinary Collaborative Design Environment Using Agents In A Virtual World', *Automation in Construction* vol. 16, pp. 37-44.
- Schmidt, K. and Simone, C. (1996) 'Coordination mechanisms: Towards a conceptual foundation of CSCW systems design' *Computer Supported Cooperative Work* vol. 5, no 2-3, pp. 155-200.
- Schroeder, R. (2002) *The social life of avatars: presence and interaction in shared virtual environments*. Springer, London.
- Schroeder, R. and Axelsson, A.-S. (2006) *Avatars at work and play: collaboration and interaction in shared virtual environments*, Springer, London.
- Schroeder, R., Heldal, I. and Tromp, J. (2006) 'The Usability of Collaborative Virtual Environments and Methods for the Analysis of Interaction' *Presence: Teleoperators & Virtual Environments*, vol. 15, no. 6, pp. 655-667.
- van Nederveen, S. (2007) 'Collaborative Design In Second Life', in *Proceedings of the Second International Conference World of Construction Project Management*, Netherlands, 2007.
- Wadley, G., Gibbs, M. and Benda, P. (2007) 'Speaking in character: using voice-over-IP to communicate within MMORPGs', in *Proceedings of the Fourth Australasian Conference on Interactive Entertainment*, Melbourne, 2007.
- Williams, D., Caplan, S. and Xiong, L. (2007) 'Can you hear me now? The impact of voice in an online gaming community'. *Human Communication Research*, vol. 33, no. 4 pp. 397-535.
- Yee, N., Bailenson, J. N., Urbanek, M., Chang, F. and Merget, D. (2007) 'The unbearable likeness of being digital: the persistence of nonverbal social norms in online virtual environments.' *Cyberpsychology and Behaviour*, vol. 10, no. 1, pp. 115-121.

Character Sharing in World of Warcraft

Nelson Wong¹, Anthony Tang², Ian Livingston¹, Carl Gutwin¹, and
Regan Mandryk¹

¹University of Saskatchewan, Canada; ²University of British Columbia, Canada
nelson.wong@usask.ca; tonyt@ece.ubc.ca; ian.livingston@usask.ca;
carl.gutwin@usask.ca; regan.mandryk@usask.ca

Abstract. Many online games are played through characters that act out players' intentions in the game world. The practice of *character sharing* – allowing others to use one's characters, or using others' – is prohibited in many RPGs, but anecdotal evidence suggests that the practice is common, and that it may play an important role in the game. To shed light on this little-known form of collaboration, we carried out a large-scale survey study to investigate character sharing in one RPG, World of Warcraft. We analyze and report on 1348 responses, providing a detailed picture of sharing practices and attitudes. We found that character sharing is common (57% of respondents reported sharing) and that sharers have a wide variety of motivations and concerns. In addition to showing how character sharing works, the study also provides new perspectives on several themes in CSCW, including conceptions of sharing, online identity, and mediating artifacts.

Introduction

In role-playing games (RPGs) players create a character in an imaginary world, acting in that world through the role of their character. The first RPGs were adventure games such as Dungeons and Dragons, played by small groups in real-world social settings. Many RPGs have now been developed for online play, commonly involving thousands of active characters in a persistent game world. These massively multiplayer online RPGs (MMORPGs) have become very popular, with millions of players worldwide (Woodcock, 2008).

MMORPGs are different from real-world RPGs because the game world is often controlled by a game publisher. Thus, players are subject to the publishers' regulations whereas real-world RPGs are governed by the players themselves. One regulation in many MMORPGs is the prohibition of *character sharing* –

where a player uses a character that belongs to another player (we consider both lenders and borrowers as sharers) – and there can be severe penalties for sharing (Blizzard, 2009). Nevertheless, anecdotal evidence suggests sharing still occurs (e.g., Jonk, 2007), indicating that it may be an important group behaviour in MMORPGs. Because of its outlaw nature, character sharing is rarely discussed openly; consequently, very little is known about this kind of collaboration.

Our goal in this paper is to shed light on this shadowy practice. We report on an investigation that used discussions with gamers and a large-scale survey to understand when, why, and how character sharing occurs in online RPGs.

The results of our study confirm that character sharing is not only common and widespread (57% of all respondents stated that they share characters in one way or another), but that it is also an important vehicle for collaborative gameplay—one that players rely on to accomplish a variety of goals. Borrowers and lenders engage in a unique type of sharing relationship, the nature of which varies based on players' attachment to their characters, their motivations for sharing, and their relationship toward the other member of the sharing relationship.

We make three main contributions. First, we uncover and document a common real-world group activity that until now has been little known and poorly understood. Second, we suggest design possibilities to better support character sharing, enabling the coordination and communication that underlie this practice. Third, we show that character sharing is a useful case study for several CSCW concepts – showing how it is a novel type of sharing, providing insight into players' relationships with their online identities, and suggesting that characters are mediating artifacts that both retain and convey experiences and state changes.

Background

Our study explores character sharing in *World of Warcraft* (WoW), an MMORPG published by Blizzard Entertainment. We set the scene by introducing relevant game concepts and terminology, and then briefly review research on WoW, online representations of players, and identity.

WoW was released in November 2004, and is the leading MMORPG game with over 11.5 million subscribers (Blizzard, 2008). Like other MMORPGs, WoW combines a predesigned story world with a character system that allows players to create narratives through in-game action and interaction (Pearce, 2004). Players create a character who is a member of one of two warring factions. Many aspects of a character can be customized, including sex, race, and clothing. The most important feature is a character's class (i.e., their job or role), which determines what skills, abilities, and equipment a character can gain and use. The differences between classes define the specific play style of a character: for example, a mage would use magic almost exclusively, whereas a warrior would generally use weapons. Characters gain experience as they are played, and with

enough experience a character attains a new level; when this happens, they are granted new skills and abilities.

A *guild* is an in-game association organized by players to accomplish in-game goals (Ducheneaut et al., 2007). One of a player's primary activities in WoW is participating in *raids* (large-scale activities involving several players) organized by these guilds. Reasons for participating in raids include searching for valuable items, and defeating hard-to-kill monsters.

A player connects to WoW using a password-protected account which is purchased and maintained with a monthly service fee. A player can have multiple characters per account. The use of this account is governed by an end-user license agreement. To enforce this agreement, Blizzard employs Game Masters (GM), whose primary job is to police in-game behaviour. In the event of a violation – such as account sharing – Blizzard may suspend or cancel the account.

The success of WoW, and its popularity among players of diverse backgrounds has made the game the subject of several research projects. Topics that have been explored include player demographics (Yee, 2006), motivations for playing (Yee, 2007), player behaviors (Yee and Bailenson, 2007), social dynamics in the game (Ducheneaut et al., 2006), gaming culture (Lindtner et al., 2008), learning in the game (Nardi et al., 2007), and collaboration (Nardi and Harris, 2006).

Our work on character sharing was also informed by studies of on-line identity and on-line representations of people. The concept of self in virtual worlds has only become common in recent years (e.g., Turkle, 1995). Research has considered how digital selves and online personas link to the virtual environment, and the nature of the relationship between people and their online identities (Donath, 1998). Previous research has shown that there is a wide range of these relationships, and that the connections between online personae and their creators are highly personal (e.g., Donath, 1998; Bessiere, 2007; Blinka, 2008).

These relationships can be affected by the nature and organization of the game genre in which the online identities exist. Role-playing games differ from other genres and from more traditional narratives in that the process of character configuration is dynamic, evolving, and determined by the players themselves (Pearce, 2004). Whereas a key factor in generating emotional responses to characters in traditional linear narratives is through empathy (Raney, 2004), interactive computer games put much more emphasis on agency, where the player controls their character and shapes the game's events (Tomlinson, 2005; Pearce, 2004). The balance between agency and empathy in RPGs may change the way players feel about their characters, and we return to this issue later in the paper.

A Survey Study of Character Sharing

Little is currently known about character sharing practices, so our initial research questions concentrated on four basic issues: *whether it happens* (what is the

prevalence of character sharing in a major online game), *why it happens* (what are players' motivations for sharing characters), *how it happens* (what are the particulars of character sharing practice), and *what factors are considered* when players decide whether or not to share a character.

To answer these questions, we designed a questionnaire to ask players of online RPGs about their character sharing practices and motivations. We developed the questionnaire through discussions with several current players, and then advertised a web-based version of the survey to WoW players.

Study Methods

We developed a web-based survey with a mixture of closed-response (check-one, check-all, and yes/no questions), short answer, and open-ended questions. The survey asked players for basic demographic information, the frequency and duration of their character sharing practice, their motivations for and reservations against sharing, and experiences with character sharing. Respondents went through one of four different paths in the questionnaire depending on whether the respondent was a borrower, a lender, both, or neither (39, 44, 69 and 12 items). Respondents spent an average of 12 minutes completing the questionnaire.

We deployed the survey for a two-week period in July 2008, and recruited participants by posting an invitation on a popular WoW forum (forums.worldofwarcraft.com). This site, frequented by both WoW players and representatives of Blizzard, is a sanctioned real-world community that allows players to ask questions and discuss in-game issues. Because respondents are WoW forum visitors, they are likely to be enthusiastic about the game and thus may not be a fully-representative sample of the general population of WoW players. However, our invitation did not mention character sharing, only stating that we were interested in studying "the playing habits of people who enjoy MMORPGs" with a link to the survey. We believe that our results are indicative of trends in the general population of WoW players.

Participants. During the two weeks that the survey was available, we received 1476 responses. We discarded 128 responses that were incomplete or from players younger than 18, leaving 1348 legitimate responses (1210 men, 112 women, 26 no response) for our subsequent data analyses. Respondents ranged in age from 18 (the minimum allowed for the survey) to 65, with a median age of 26.

Our survey attracted a wide range of participants, skewing slightly toward dedicated gamers: 62% rated themselves as 'regular' players, 24% as 'hardcore', 13% as 'casual', with 1% abstentions. We asked users to classify their player type based on descriptors adapted from Bartle's (1996) descriptions (Achiever, Explorer, Killer, Socializer). The majority of participants (52%) identified themselves as Achievers, meaning that they "set game-related goals, and vigorously set out to achieve them" (Bartle, 1996). 19% of respondents identified

themselves as Explorers, 13% as Killers, 8% as Socializers, and 8% either not responding or stating that that they did not identify with any of the categories.

Data Analysis. For check-all-that-apply questions, we solicited additional information through a free-form follow-up question. After coding the open-ended responses, we integrated the user-supplied answers with the original check-all-that-apply answers for further analyses. For each type of multiple-choice question, we present the results as percentages of the respondents who answered that specific question. The number of respondents for each question varied due to the participants’ varying paths through the survey.

Does Character Sharing Happen, with Whom, and How Often?

Our results show that character sharing is both widespread and frequent (see Figure 1). 57% of respondents stated that they shared characters in some way. Of these, 74% reported *lending* characters to others, while 94% reported *borrowing* characters from others. Of the 43% of respondents who do not share characters, 84% of these report having made an explicit decision not to share, while the remaining 16% report not having had the opportunity to share.




Figure 1. Sharing among our participants.

We asked survey participants with whom they decided to share, and how long they allowed the sharing arrangements to go on. The four main types of people that participants reported sharing with were family, real-life friends, in-game friends, and fellow guild members (see Table I).

	Family	Real-life friend	Game friend	Guild
Loaned a character to...	20%	50%	27%	3%
Borrowed a character from...	13%	37%	36%	14%

Table I. Sharing percentages with different types of people.

The only major difference in lending and borrowing patterns is in sharing with guild members. People are willing to borrow from these people, but less likely to lend; this may be because guild relationships are not as strong as personal relationships, but may also arise because of so-called ‘guild accounts’ where all guild members can access the guild account’s characters.

Participants reported two main types of sharing arrangements. The most common was ‘one-time’ sharing where the borrower used the character once for a particular purpose (40% of total sharing). In these cases, borrowers were expected not to log in again afterwards. Several lenders reported temporarily changing their password for the duration of the share, and then changing it back afterwards.

The second type of sharing arrangement was longer term, and allowed the borrower to repeatedly log into the account (25% of total sharing). In some cases

this arrangement was used because the in-game task was time-consuming (e.g., obtaining several copies of hard-to-get items); in other cases players had long-standing agreements with friends or their guild that characters could be used when needed (e.g., where a guild “[has] access to our main warrior's account”).

In a few cases, there were mutual long-term arrangements within a group. Participants described situations in which all player accounts were known to the entire group, and where players were welcome to use others’ characters at any time to achieve the goals of the group. In one case, it appeared that these accounts did not even have real owners, and were instead owned by the entire group:

[A guild] I belonged to had a "shared guild account." This account was given from a player who stopped playing to the owner of the guild. This account information was then given to all trusted members... to use the characters if it was needed.

We also asked participants how many times they had shared characters. For those who reported lending, people had lent characters an average of 10.8 times; borrowers reported that they had borrowed characters 9.1 times on average.

Motivations for Sharing – Why do People Share Characters?

In this section, we examine motivations for character sharing, illustrating that character sharing is largely motivated by a desire to experience the game more fully. Participants identified 22 reasons for sharing characters, but four groups of those made up 65% of all responses (described in the following sub-sections). The ten most frequent reasons for sharing are presented in Figure 2.

Sharing to experience new things (72% of sharers)

Each character in WoW experiences the game in different ways: for instance, each of the two warring factions has a unique story unavailable to the other faction. Most players advance through the game using a single character, and few invest much time in alternate characters; consequently, most players only experience




Figure 2. The main motivations for borrowing and lending characters.

gameplay through a single character. Our survey shows that many players are curious about other aspects of the game and other character classes – especially high-level characters who gain access to special content – and how those character classes experience the game. In our sample, many respondents reported sharing characters to play different characters (58% of lenders and 47% of borrowers) and to experience different aspects of the game (58% and 44%), for example:

Sometimes [I borrow] to try a class that I haven't played before, and that I am interested in leveling, but don't want to max out to find out that I don't like it.

Similarly, several respondents loaned characters to real-life friends to allow them to try the game: 55% of lenders and 20% of borrowers reported having shared characters for this reason. In these situations, the benefit is primarily to the borrower who is able to have a different (or new) play experience.

Together, these three motivations (to play different characters, experience different aspects of the game, and try the game) were reported by 72% of sharers.

Sharing to ensure adequate resources for a raid (43% of sharers)

A *raid* in WoW is an organized group activity where a team of players attempt to achieve an in-game objective defined by the game designers (e.g., defeat a monster). Raiding parties contain 6-40 characters, with each character typically playing a specific role (e.g., *damage dealers* who attack the enemy, *healers* who restore other characters' health). Coordinating the many players needed for a raid is often difficult: owners of some important characters may not be available at the scheduled raid time. In these situations, it is common to loan important characters to a player who is available for the raid. For example, as one participant stated:

He asked me to play his account as we were sho[r]t a healer and he couldn't make it that night.

29% of lenders and 40% of borrowers reported sharing characters for this reason (to access unique skills of the shared characters). Sharing benefits the raiding party because the group needs the skills of the shared character: often, raids cannot be carried out without the appropriate balance of roles. The owner of the shared character also benefits, because their character receives a share of the spoils from the raid. Lending for raids is most often a short-term arrangement (lasting as long as the raid); however, this situation was also a reason to set up a more permanent lending arrangement. For example, one respondent stated:

sometimes we need a warrior to tank a boss but we don't have a warrior online, but we have access to our main warrior's account, [so] I'd log on the warrior and bring him to the fight then after the fight go back to playing my character.

In addition, a few respondents reported having a 'guild account' (as discussed above) that is accessible by all guild members and that was used for raids.

Sharing to advance a character (38% of sharers)

Leveling is the activity of moving a character to a new experience level and often involves the completion of dull, repetitive tasks. Although these tasks are part of

the game, players often consider aspects of leveling a necessary evil. To reduce the effort and pain of leveling, some players lend their character to a friend or even to a private business that will carry out some of the required tasks. This type of sharing is different from other reported types, as it primarily benefits the lender rather than the borrower. In our survey, 20% of lenders and 33% of borrowers reported sharing characters to level a character more quickly. Although most sharing in these situations was intended to avoid repetitive work, some cases involved a sincere interest in helping another person – for example:

My good friend has trouble leveling her characters, and not being able to participate alongside her friends and her husband because she was too low-level; [this] was very distressing to her, so I helped her out.

Sharing for leveling is more controversial among players than other reasons for sharing. Many people saw it as cheating, since the character was no longer a true reflection of the owner's skill (e.g., "playing a character that's been leveled [by someone else] feels like cheating"). It was regarded in the same light as allowing a character to be advanced by a 'bot', a practice that is also disallowed.

Sharing to learn new techniques for playing the game (33% of sharers)

The WoW user interface is highly customizable, allowing players to modify and tailor in-game commands to their specific needs; for instance, macros may be recorded to automate sequences of commands. However, in-game tricks or techniques are often difficult to explain to newcomers. Sharing a character allows the borrower to learn these enhancements – in these cases, it is not so much the character that is shared as much as the customized environment.

Many respondents reported employing character sharing to either teach another player about some aspect of the game (e.g., instances where certain macros are useful), or to learn from another person. Often this type of sharing was carried out in a co-present environment, so that the lender and borrower could more easily talk about the interface. In our sample, 20% of lenders and 26% of borrowers reported sharing characters for this reason, showing that customization – and community support for customization through sharing (Mackay, 1990) – has become common in WoW.

Details of Sharing Practice – How Does Character Sharing Occur?

This section looks at the details of character sharing: setting up the arrangement, coordinating the use of the character, and finding out what happened afterwards.

Managing the handover: transfer and scheduling

Accounts in WoW are protected by a username and password, and so the actual transfer of a character involves the transfer of account details. This information is typically sent through email or IM (85% of lenders) or by logging in and letting a co-present borrower use the account (31%). The more complex handover issue,

however, is that of scheduling to avoid conflict on the account, because if another player attempts to log in to the account while the first login is active, then the first player will be disconnected, or *kicked*. Beyond being an inconvenience, this can also cause serious problems if the character is in the middle some important activities. For example, one participant stated:

[I] once logged on to my character while a friend was using him [...] the character was underwater when it happened and the delay in transition caused him to drown.

Because only one person can be logged in to the account at once, organizing and following a schedule is crucial. Respondents relied both on large-scale coordination (e.g., “I only loan my characters to others when I'm not playing the specific game at that time”) and finer-grained scheduling (e.g., “I told the person they could use [my characters] while I was at work so between the hours of 9-5”).

Respondents also felt that multiple logins could draw the attention of the Blizzard game masters, which could result in banning of the account. Consequently, most borrowers (78%) indicated that it was important for the borrower to inform the lender before logging in as the shared character.

Limiting the borrower: rules and restrictions

Most of the lenders in the survey (74%) placed restrictions on how shared characters could be used. Respondents stated many different rules that were based on the specifics of characters, situations, and the borrower themselves. The most common restriction (mentioned by 44%) relates to the use of a character's in-game resources such as money and items, because they may be difficult to reacquire. For example, a common set of rules were:

don't sell/delete anything without asking. Don't use crafting materials without asking. Don't re-spec [(change character attributes)] unless I ask you to.

Another common rule was similar – during gameplay, irreversible decisions occasionally need to be made (e.g., selling unique items); consequently, many lenders stated that they tell borrowers to avoid making such decisions, or only lend to other players who ‘already know what not to do’ with the character.

Getting the character back: finding out what happened

Characters are ‘returned’ either implicitly through the scheduling arrangement, or by the borrower notifying the lender that they are finished. This is not, however, the end of the sharing lifecycle: after the character is returned, the majority of lenders (67%) also want to know ‘what happened’.

Interest was highest in the outcome (40% of lenders) – the success of the borrower's task, in-game tasks that had been accomplished, changes in the character's inventory, and the character's game world location. Lenders gathered this information in two main ways. First, they spoke with the borrower, either by voice or online (several respondents stating that a real-time medium was necessary to allow clarifying questions to be asked). Second, lenders also gathered information by inspecting their characters: 42% of lenders reported studying the

character's item inventory to determine which items (e.g., gold or equipment) had been used, obtained, or sold. The inventory functions as a persistent, indirect record of activity: for example, it can show that a character has been in battle (e.g., health potions depleted), or has succeeded in a task (e.g., new items acquired). In addition, lenders also checked the inventory to ensure that the borrower had not wasted or given away items – one participant reported that he went so far as to take a screenshot of the inventory before lending a character, and then checked the screenshot against the character's inventory afterwards.

Some lenders were also interested in other experiences that did not result in changes to the character, although this was mentioned less frequently (27% of lenders). People stated that they were also interested in the actual experiences that the character had while 'away': when the character was played, what monsters they fought, how items were obtained, and whom they encountered in the game. For example, one lender wanted "to know who in the game my character has encountered so I am not confused later."

Factors in Deciding Whether or Not to Lend

We asked sharers a check-all-that-apply question about their concerns when lending and borrowing characters. We also asked non-sharers their reasons for not sharing characters. Results are presented in Figure 3, and below we detail the five most frequent reasons.

Fear of being caught

Character sharing requires account sharing, which is against the publisher's terms of agreement for the game. The fear of being caught and punished is a major concern for players, and a serious deterrent for those who choose not to share (it was indicated by 57% of non-sharers and 37% of sharers). As one person stated,

[playing someone else's character] can be really tense. It feels strange playing on someone else's account and knowing that you're breaking the ToS

Blizzard watches for infractions such as account sharing, and users mentioned issues with logins from distant IP addresses or multiple logins. For example,

I know a few people who got banned because a GM [(game master)] noticed weird login / IP addresses on their accounts.

Identity

Respondents stated that three kinds of identity issues were important. First, some players identify strongly with their characters and consider them to be extensions of their selves (38% of non-sharers felt that characters were a reflection of one's personal identity, and 22% of sharers also indicated this response). This strong relationship to on-line avatars has been reported numerous times in past research (e.g., Turkle, 1995; Blinka, 2008), and for many players, this was the primary reason for not sharing:

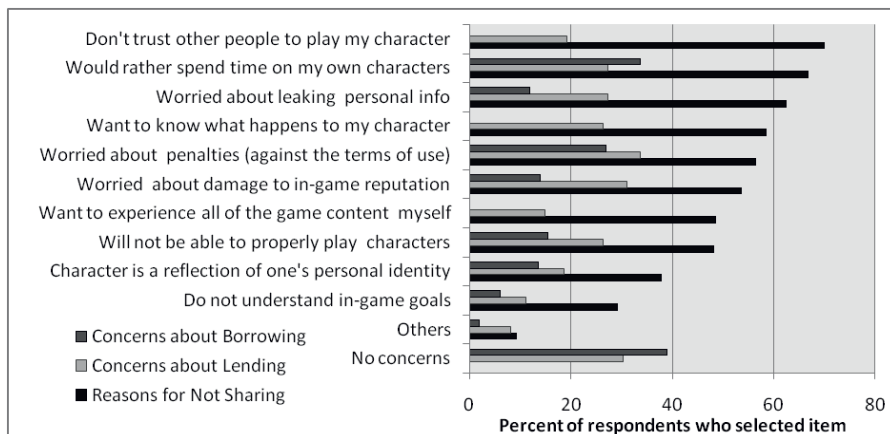


Figure 3. Concerns about borrowing and lending, and reasons for not sharing.

I feel my characters are a personal incarnation. The personality that they are is me and people come to know this and enjoy being around me due to this. When someone else plays my characters I feel it throws things off in a way.

Second, even if they did not see characters *as* themselves, many respondents felt that their characters *stood for* their real-world identity, reputation, and social standing. For example, one person said that “a character is an icon of one's social identity in the online world;” another stated that “a character is a reflection of my personal identity.” Players who felt this way were sometimes willing to lend their characters, but were concerned about how the borrower would play the character (e.g., one stated “I don’t want my reputation to be ruined”). Accordingly, many borrowers reported playing a shared character with greater care (so as to not damage the lender’s social standing).

These responses suggest that in some interactions there is a clear separation between the character and the real-world person behind it. A third identity issue that is strongly related involves the practical realities of carrying on real-world interactions through in-game characters. Many respondents mentioned problems arising from the fact that during sharing, a different real-world person is now behind the character. These cases of mistaken identity can lead to confusion and out-of-context communication. In some cases, mistakes lead to social *faux pas*:

The [owner] used to log in at a different time than me and chat with others... and became very friendly with someone else. Needless to say... the conversation that came my way when I happen to log in on a day off from work was not something I was expecting... especially since the friend using the account was a she and I am a he. It was rather embarrassing for all concerned.

Borrowers mentioned several times that this issue leads them to avoid starting conversations when playing another person’s character: as stated by one person, “my biggest concern is their in-game friends talking to me, I'm not familiar with them so I don't know how to respond to them.” Problems caused by mistaken identity led several borrowers to consistently reveal who they were (i.e., not the owner) when others engaged them in conversation. Most borrowers (54%)

indicated that it was appropriate to inform others in this way, but lenders were evenly split as to whether borrowers should do so.

Characters as Investments

Another factor that lenders consider is the value of the character, and the potential loss that could occur if something goes wrong. Advancing a character through WoW, and obtaining gold and equipment, requires a considerable investment of time; many respondents stated that they thought more carefully about sharing higher-level or wealthier characters, and imposed rules about how borrowers should act (as described above). For example, a lender stated:

[I] have a huge amount of gold and items...I don't like the feeling of my friends, even my best friend, playing on my characters and not knowing exactly what they did when they played my characters.

In addition, the idea of characters as investments was raised as a concern for borrowers; that is, that playing a shared character would be a waste of time since the value would go to someone else (34% of borrowers considered this a drawback). The idea of character as possession (rather than as persona) warrants further investigation and we return to this idea in the discussion.

Trust and Security

Trust in the borrower was a major concern for players: 70% of non-sharers stated that this was a factor in their decision, as did 19% of lenders. Sharing relationships generally follow real-world trust patterns – as shown in Table I, characters are lent primarily to friends and family members. Both non-sharers and lenders are concerned about whether they can trust the borrower to protect their reputation (54% of non-sharers, and 31% of lenders), and to play the character properly (48% and 26%). Even maintaining interface settings is a concern:

I spent maybe an hour going over screenshots in an attempt to re-create my UI toolbars after that incident.

Thirty percent of lenders, however, reported no concerns with sharing their characters, suggesting that a sizeable minority of lenders either do not mind what happens, or that there is implicit trust, as stated by one participant who said,

honestly, I don't care. Unless of course its something serious, but I wouldn't expect anything like that to happen.

In addition, many players perceived character sharing as a potential security risk: 62% of non-sharers and 28% of lenders stated that they were concerned about personal information when sharing characters. Security problems can occur in several ways: first, the account contains considerable personal information that could be given out or lost; second, if a borrower changes the account's password, a lender could lose the account completely. These concerns led to practices such as changing account passwords every time a character is shared (as described above). Last, players were concerned about risks from the borrower's computer:

I don't give my account information away, because [although] I trust friends not to mess with my characters, I do not know if they protect their computer against hackers. I want my account to be safe.

Summary of Survey Results

Our survey provides evidence about the existence, prevalence, and complexity of character sharing in World of Warcraft. In summary:

- *Sharing is frequent and widespread.* The majority of respondents have lent or borrowed characters, and have done so many times.
- *Sharing has two main patterns:* one-time sharing, where characters are returned once a particular task is completed, and longer-term repeated sharing.
- *Sharing is used for several purposes.* There are many different reasons for sharing characters, the majority of which are not considered to be cheating.
- *There are several types of player-character relationships.* Players indicated that they think of their characters in many ways: extensions of themselves, as valued possessions, and even as throwaway objects.
- *Identity is a main concern.* Online identity issues are a major factor in sharing, leading some people to avoid sharing, and others to be careful about protecting their reputations and avoiding problems with mistaken identities.
- *Change awareness is important.* The majority of lenders want to know what happened to shared characters, and use both in-game (e.g., character inventory) and non-game channels (e.g., telephone), to obtain this information.
- *Communication about sharing is required.* The practical details of sharing involve considerable communication – for transferring account information scheduling, setting rules, and reporting what happened to the characters.
- *Sharing is not well supported.* The lack of any in-game support for character sharing forces people to engage in risky practices and to use tools and mechanisms (such as screen shots for awareness) that are often awkward.

Discussion

Our study reveals many of the details of character sharing, a collaborative practice that has not been studied before in CSCW. However, the broader value of our study is that character sharing raises new questions for a number of existing CSCW topics – in the next sections, we discuss the ways that character sharing may be able to shed light on research into sharing, on issues of player-character identity, and on characters as a mediating artifact in the articulation work of sharing. In addition, we consider the question of whether character sharing should be better supported by game companies, and present several design ideas that could help to provide this support.

Character sharing is a different kind of sharing

There are fundamental differences between the sharing of game characters and the types of sharing that have been studied previously in CSCW, including program customization files (e.g., Mackay, 1990), shared folders (Voida et al., 2006), music sharing (e.g., Brown et al., 2001), and photo sharing (e.g., Miller and Edwards, 2007). The main difference is that sharing of files, music, and photos involves digital objects that can be trivially and transparently copied, meaning that people are actually sharing a copy of the artifact rather than the owner's original. In contrast, characters in on-line games are unique and cannot be copied, since they are tied to the owner's unique account with the game publisher.

This means that sharing practices and people's attitudes toward the shared object are dramatically different. With music or file sharing, there is no concern about getting the shared object back again, and the idea of sharing, in part, implies the idea of making the artifact public (particularly with photo sharing). With copy-based sharing, there is also no need to maintain awareness of what happens with the shared object while in the borrower's possession. Although the 'lender' may still take an interest in what the borrower does with the object (e.g., makes a new version of a song or adds to a customization file), the original version is still in the owner's possession, and lending creates a version tree rather than accumulating changes to the original object itself, as occurs with a WoW character.

The fact that there is only one copy of a WoW character means that character sharing is more like sharing real-world objects like cars or bicycles than it is like sharing other types of digital objects. In particular, the owner sees real value in the actual object being shared, and so considerably more thought must be given to decisions about when and with whom to share. Thus, we see many comments about whether the lender can trust the borrower to use the character appropriately – concerns that generally do not occur in copy-based sharing. Player comments about this issue sound very similar to what goes through one's mind before lending a valued real-world possession, such as a car or a book, to another person (e.g., as one participant in our survey said, "I would want to know whenever someone wants to use my car, the same goes for my character").

There has been very little CSCW research done on this type of sharing; work exists in areas such as deception in Usenet discussions (Donath, 1998) and group computer accounts (Egelman et al., 2008; Muller and Gruen, 2005), but there is much that could be done in this area. For example, an issue raised by our study was the wide range of value that lenders placed on their characters – from treasured possessions that would never be lent out, to throwaway objects with little value. Part of the reason for this wide range is that the actual creation of characters is easy, and so the value of a character does not arise only from its mere existence (as it would with some kinds of physical objects). Instead, it appears that value is primarily created by the degree of the owner's involvement in the character (e.g., the investment of time and effort to reach a particular level).

Therefore, characters are ‘self-built,’ somewhat like handmade furniture or pottery, and character sharing shows similarities to sharing these types of personally-meaningful items.

In a different way, however, character sharing is similar to other types of digital sharing – these types of group activity are interesting for CSCW in that they raise the question of where in a sharing relationship the collaboration actually occurs. Character sharing appears to be a type of articulation work, in that it enables some other end goal; but only in some sharing arrangements (e.g., using a character for a raid, or working towards a level) does there appear to be a common goal between the lender and the borrower. In other cases, such as allowing others to try out the game or try out a different type of character, there does not seem to be a clear group goal – in that sharing allows one person to have an individual experience that they could not otherwise have. Character sharing is therefore a mechanism for social interaction in the larger community (and in this domain, helping others to new experiences could indeed be part of the larger shared goal), as much as it is a coordination mechanism for ‘getting things done,’ and thus contributes both to thinking about focused work activity and to research on the broader social issues that have been considered in other studies of digital sharing (e.g., Brown et al., 2001; Håkansson et al., 2007).

Characters as a new kind of mediating artifact

The artifacts that are transferred between people in collaboration can store and show information that aids articulation work. As stated by Schmidt and Simone (1996), the artifact “mediates articulation work as well in the sense that the artifact acts as an intermediary between actors that conveys information about state changes to the protocol under execution” (p. 179). It is clear that characters in WoW play this role of mediating artifact – for example, in situations where lenders inspect the character’s inventory to determine what items have changed.

Character sharing extends this idea, however, in that characters not only show state changes that have occurred during the share, but also ‘contain’ the in-game events and happenings that the character has experienced. These experiences are often as important to lenders as are changes to gold or equipment, and several people stated that they were reluctant to lend characters because they didn’t want to miss out on what happened. Thus, the story of the changes is often as important as the changes themselves, and characters can be seen as mediators of experiences as well as representations of the state of the sharing arrangement.

There is currently no way to extract these experiences from the character, however. Although research into edit wear and read wear (Hill et al., 1992) has considered the idea of recording and displaying a wide variety of interaction history (and these techniques could also benefit character sharing), prior work has not considered the artifact’s own experience (e.g., the character’s adventures

rather than its state changes) as history that could be recorded. The fact that characters' experiences are understandable to players means that there are new opportunities for characters to relate and share this information – such as the possibility of asking the character questions about their adventures, or the possibility of playing back experiences (as discussed below).

Character sharing exposes identity issues in online environments

We were surprised by the degree of willingness to share characters, and by the almost casual attitude towards sharing that we saw from some of our respondents. Although we did not focus specifically on identity issues, answers and comments suggest that there may be several ways that character sharing can illuminate the ongoing interest in the relationship between players and their in-game avatars. In particular, our study suggests that WoW has a wider range of relationships than have been reported before, and that some of these involve less of an identity connection between player and character – sometimes to the point where a character is as much a possession as it is a persona.

The way in which players see their relationship with their characters contributes to their attitudes towards character sharing, and through our participants' comments, we saw several different types of relationship. As discussed earlier, some players strongly identified with their characters and thought of them as extensions of their own identities and personalities. Other players considered characters as a symbolic representation of their real-world self. At the other end of the spectrum, several players talked about characters as objects quite separate from themselves; as discussed above, several people thought of their characters as property and explicitly referred to them as such (e.g., one participant stated that “a character is a personal possession;” another said “objectively, a character is something that is property”).

The degree to which a player identifies their characters with themselves could have substantial effects on sharing practices. The more closely a player associates themselves with a character, the less likely they would be to lend it out, and the more concerned they would be about the character's behaviour while in someone else's care (e.g., one lender told borrowers “don't act like an ass”; other lenders stated rules for the borrowers, as described above). In contrast, seeing characters as possessions could lead to much more willingness to share, and more interest in the character's inventory as opposed to their behaviour (as one lender said “no rules, I don't really care that much [...] I would prefer them not to delete or sell my stuff”). This difference could also explain people's different opinions on whether lending characters is cheating – at the one extreme, obtaining skills or materials without personal investment would be similar to the falseness of getting cosmetic surgery; at the other, it would be no more devious than letting someone else tune up your car or fix your bicycle.

The ways in which characters are created and manipulated in virtual worlds may have an effect on how much players identify with the character. For example, the player's avatar in *Grand Theft Auto* has a pre-determined name and backstory, and although his appearance can be customized, it is unlikely that players see this character as a representation of their actual selves (except in a vicarious sense). Other games provide different creation mechanisms that can allow a closer bond between player and character; but the details of character creation and management still affect the relationship. One way that World of Warcraft differs from other virtual worlds is that an account can contain multiple characters; in contrast, worlds like Second Life allow only one avatar per account. The ability to create multiple characters may be important for identity issues because it makes a clear break from the character-equals-player relationship, and makes possible the existence of multiple characters in which the player has not invested time.

Few of the participants seemed to think of their characters as only possession, but the language people used to talk about their characters (e.g., comparisons to cars or other objects) showed that the idea of characters as property is present at least for some players. This idea has not been widely considered in CSCW research before, and presents several new opportunities for further research – for example, arguments about cheating in virtual worlds could be informed by an understanding of this identity issue.

Should character sharing be supported?

Our survey shows that character sharing is already widespread, and that not all aspects of this practice are likely to be harmful to the game publishers or the in-game experience. In addition, there are many benefits in sharing – it brings people to the game, it helps people get greater enjoyment out of the gameplay, and it aids the development and maintenance of social groups both in game (guilds) and in the real world (local and broader communities of players).

Our findings suggest that game publishers could benefit from thinking about ways to support different aspects of character sharing. Although there are several issues at play in this debate, one of particular interest to CSCW is the question of whether (and how) companies can support types of collaboration that enable prohibited activities. Discussions of articulation work in CSCW have often highlighted the failure of groupware systems to support the essential activities that go on behind the scenes (e.g., Schmidt and Simone, 1996). World of Warcraft can be seen in exactly this light – as a groupware system that fails to recognize the behind-the-scenes work (i.e., character sharing) that is needed to accomplish a variety of tasks and aims in the game and in the larger community of players. However, it is not the case that WoW's designers have simply failed to notice an important aspect of group work; the problem is that the activities are prohibited. This poses the question of whether an activity should be supported when doing so

makes it easier to engage in disallowed behaviour – game publishers may believe that doing so would be seen as legitimizing these activities (Birnholtz et al, 2008). In addition, there is the possibility that legitimizing character sharing could dramatically change the way characters are developed and used (e.g., open rental of characters or more widespread sharing beyond a player’s immediate social circle), and raises many questions for further study.

Although we do not expect Blizzard to embrace character sharing in the near future, there are possible ways forward that could obtain some of the benefits of character sharing without compromising account security, and without ruining the experience for other players. In addition, it would be relatively simple to sanction and support certain aspects of the practice if players are willing to live with the (mild) cheating that it allows. In the next section, we consider some of the ways that character sharing could be better supported, if there was a willingness to do so either in WoW or in some new role-playing game.

How could character sharing be better supported?

The issues and attitudes shown in the survey suggest several design ideas that could provide more explicit support for different aspects of character sharing. In the following paragraphs we describe seven design changes that were closely aligned with the results of the survey, and that could be feasibly implemented.

Decoupling accounts and characters. The security risks of current character-sharing practices could be dramatically reduced by allowing characters to be played from different accounts. Each gamer would still need an individual account, but the characters would no longer be tied exclusively to it.

Different levels of access. The owner of a character should be able to control what a borrower can do to and with the character. This could be done by locking certain functions of a character to borrowers (e.g., item usage).

Tracking changes. Systems should provide change-awareness information to lenders when characters are returned. This could be done by simple visualizations, such as highlighting the changes in the inventory.

Playback tools. Another way to support change awareness is to provide playback tools (e.g., videos or screenshots). In addition to awareness, playback gives lenders a way to participate in experiences that they have missed.

Private sticky notes. Characters could be used as a repository for asynchronous communication between borrowers and lenders. Lenders could, for example, attach sticky notes to a character to tell borrowers what they should do next with the character, and borrowers could use them to report what has happened.

Spectator mode. Spectator mode would allow players to observe gameplay through another player’s view. Such a mode in WoW could reduce unwanted ‘kicking’ of a borrower, and could also provide real-time feedback. Sharers could even trade control of the character for a collaborative gaming experience.

Identity indicators. Knowing who is controlling the character is important. Graphical indicators, such as a halo around the character, could show whether the player is the owner of the character or a borrower (or even the name of the player). Identity indicators would reduce confusion and cases of mistaken identity.

These new tools and techniques could dramatically simplify practices that are currently carried out through clumsy and insecure mechanisms.

Conclusions and Future Work

Until now, little has been known about the prohibited practice of character sharing in online role-playing games. We surveyed 1348 WoW players to investigate this practice, and although this is a relatively small sample, the survey is the first to report on this shadowy form of collaboration – we show that it is widespread, frequent, and plays an important role both for in-game collaborative activities and for interaction in the larger community of players. In addition, character sharing sheds new light on several themes in CSCW: characters are a different kind of object than has been considered in studies of sharing; character sharing shows new perspectives on the relationship between a player and their online identity; and characters can be a novel type of mediating artifact that contains experiences in addition to state changes. Our work in this area will continue in two directions: first, we plan to confirm our findings through discussions with players of other MMORPGs; and second, we will further explore the issues of sharing, identity, and mediating artifacts that have been raised by our study.

References

- Bartle, R. (1996): Hearts, clubs, diamonds, spades: Players who suit MUDs. *Journal of Online Environments*, 1(1).
- Bessiere, K., Seay, F., and Kiesler, S. (2007): The Ideal Elf: Identity Exploration in World of Warcraft. *Cyberpsychology and behavior*, 10 (4), pp.530-535.
- Birnholtz, J., Gutwin, C., and Hawkey, K. (2007): Privacy in the Open: How Attention Mediates Awareness and Privacy in Open-Plan Offices, In *Proceedings of Group 2007*, pp.51-60.
- Blinka, L. (2008): The relationship of players to their avatars in MMORPGs: differences between adolescents, emerging adults and adults. *Cyberpsychology: Journal of psychosocial research on cyberspace*. 2(1), pp.1-9.
- Blizzard. (2008): World Of Warcraft Subscriber Base Reaches 11.5 Million Worldwide. <http://www.blizzard.com/us/press/081121.html>
- Blizzard. (2009): Account Sharing/Third-Party Character Advancement ("Power Leveling"). <http://us.blizzard.com/support/article.xml?articleId=21506>
- Brown, B., Sellen, A.J., and Geelhood, E. (2001): Music Sharing as a Computer Supported Collaborative Application. In *Proceedings of ECSCW '01*, pp.179-198.
- Donath, J (1998): Identity and deception in the virtual community, in Kollock P and Smith M (Eds), *Communities in Cyberspace*, Routledge, London, pp.29-59

- Ducheneaut, N., Yee, N., Nickell, E., and Moore, R.J. (2006): Alone Together? Exploring the Social Dynamics of Massively Multiplayer Games. In *Proc. of CHI 2006*, pp.407-416.
- Ducheneaut, N., Yee, N., Nickell, E., and Moore, R.J. (2007): The life and death of online gaming communities: a look at guilds in World of Warcraft. In *Proc. of CHI 2007*, pp.839-848.
- Egelman, S., Brush, A. J., and Inkpen, K. M. (2008): Family accounts: A new paradigm for user accounts within the home environment. In *Proceedings of CSCW'08*, pp.669-678.
- Håkansson, M., Rost, M., and Holmquist, L. E. (2007): Gifts from Friends and Strangers: A Study of Mobile Music Sharing. In *Proceedings of ECSCW'07*, pp.311-330.
- Hill, W., Hollan, J., Wroblewski, D., and McCandless, T. (1992): Edit wear and read wear. In *Proceedings CHI 1992*, pp.3-9.
- Jonk, J. (2007): Blizzard bans 12 wow PVP accounts. www.esreality.com/?a=post&id=1386800
- Lindtner, S., Nardi, B., Wang, Y., Mainwaring, S., and Jing, H. (2008): A hybrid cultural ecology: world of warcraft in China. In *Proceedings of CSCW 2008*, pp.371-382.
- Mackay, W. E. (1990): Patterns of sharing customizable software. In *Proceedings of CSCW'90*, pp.209-221.
- Miller, A. and Edwards, W.K. (2007): Give and Take: A Study of Consumer Photo-Sharing Culture and Practice. In *Proceedings of CHI 2007*, pp.347-356.
- Muller, M., and Gruen, D. (2005): Working together inside an mailbox, In *Proceedings of ECSCW 2005*, pp.103-122.
- Nardi, B. and Harris, J. (2006): Strangers and Friends: Collaborative Play in World of Warcraft. In *Proceedings of CSCW'06*. pp.149-158.
- Nardi, B., Ly, S., and Harris, J. (2007): Learning Conversations in World of Warcraft. In *Proceedings of HICSS 2007*, pp.79a.
- Pearce, C. (2004): Towards a game theory of game. In *First person: new media as story, performance, and game*. Ed. Noah Wardrip-Fruin & Pat Harrigan. Cambridge: MIT Press.
- Raney, A. (2004): Expanding disposition theory: Reconsidering character liking, moral evaluations, and enjoyment. *Communication Theory* 14(4), pp.348-369.
- Schmidt, K., and Simone, C. (1996): Coordination mechanisms: Towards a conceptual foundation of CSCW systems design, *CSCW*, 5(2/3): pp.155-200.
- Tomlinson, B. (2005): From linear to interactive animation: how autonomous characters change the process and product of animating. In *Computers in Entertainment*, 3(1), 5-5.
- Turkle, S. (1995): *Life on the screen: Identity in the age of the Internet*. NY: Simon & Schuster.
- Voida, S., Edwards, W., Newman, M. W., Grinter, R. E., and Ducheneaut, N. (2006): Share and share alike: exploring the user interface affordances of file sharing. In *Proceedings of CHI 2006*, pp.221-230.
- Wachowski, E. (2007): Rogue with Twin Blades of Azzinoth sells account for almost \$10,000. <http://www.wowinsider.com/2007/09/16/rogue-with-twin-blades-of-azzinoth-sells-account-for-almost-10>
- Woodcock, B. (2008): MMOG Chart. <http://www.mmogchart.com>
- Yee, N. (2006): The Demographics, Motivations and Derived Experiences of Users of Massively-Multiuser Online Graphical Environments. In *PRESENCE: Teleoperators and Virtual Environments*, 15, pp.309-329.
- Yee, N. (2007): Motivations of Play in Online Games. In *Journal of CyberPsychology and Behavior*, 9, pp.772-775.
- Yee, N. and Bailenson, J.N. (2007): The Proteus Effect: The Effect of Transformed Self-Representation on Behavior. In *Human Communication Research*, 33(3), pp.271-290.

High-Octane Work: The oil and gas workplace

Clint Heyer

ABB Strategic R&D for Oil, Gas and Petrochemicals
clint.heyer@no.abb.com

Abstract. This paper introduces the oil and gas workplace context and describes work practices observed at a large Norwegian gas refinery. Ethnographic fieldwork was carried out over a ten day period, consisting of observational studies and informal interviews. They are a small, inter-disciplinary group who are highly mobile and work in a hazardous, critical environment where mistakes can pose risk to health, safety and the environment as well as significant financial loss. Two main shift roles, field operator and central control room operator, are discussed and related to the wider workplace. Even in this technologically-advanced workplace, non-digital informational artifacts are important, often serving as bridges to support flowing activity between communities of practice and the physical and digital. Spending time in the physical plant was seen as an important way to develop an understanding of the process and to gain insight not available through a control system. The primary contribution of this paper is the detailing and discussion of an oil and gas workplace from a CSCW perspective, a context not well established in the literature, yet one that poses an interesting range of design challenges.

Introduction

Oil and gas (O&G) workplaces present manifold challenges for design and seem a promising area of future research. O&G operators are mobile, work in inter-disciplinary teams, in and in-between harsh outdoor and benign indoor environments. They make use of spectrum of tools, from complex ERP (Enterprise Resource Planning) information systems through to rudimentary mechanical tools. Additionally, the refinery or process they attend to is spread over a large space which is both logically and physically complex. High-pressure pipes, extremely low and high temperatures and explosive materials mean the

workplace is also a hazardous one, posing risks to human health and safety as well as the environment (HSE). As such, work practices, protocols and equipment are first and foremost designed to promote safety and decrease risk. This naturally impacts what kind of technology can be deployed at a site, for example a normal camera cannot be used without special precautions due to the risk of explosion. Work is of a critical nature: mistakes can cause HSE issues and/or disrupt production, both of which can have a very large financial impact.

Kvasir¹, the studied site, is a refinery situated in a remote area of Norway, supplying a large proportion of the European market and one of many onshore and offshore facilities operated by its parent company. Like most onshore facilities, it is separated into two areas, administration and plant. The three-storey administration building is where regular day shift workers have their offices and also houses the control room, meeting rooms and cafeteria. The plant is where the refining takes place, thus the location of highest HSE risk. It is located 2km from the administration building with workers transiting by vehicle or bicycle. Shift teams' priority is to ensure safe, 24-hour production, while engineers and maintenance staff focus on longer-term upgrades and repair work.

Control rooms featured in CSCW literature have frequently been transport-related, such as air (Bentley, Hughes, Randall et al., 1992), subway (Heath and Luff, 1992) and ambulance (Martin, Bowers and Wastell, 1997). We suggest that the O&G control room is different in two ways. Firstly, there is a significant amount of collaboration and interaction in the O&G control room between people of different disciplines. For example, engineers and field operators might meet in the control room to work through a fault with a control room operator. Furthermore, O&G control room operators have frequent and direct contact with the people they issue directives to (field operators) and work with them in a highly collaborative fashion remotely or co-located in the control room. Secondly, the high degree of automation in O&G facilities mean 'control work' is usually only conducted when something is amiss or when maintenance or repairs are being carried out. As such, control room operators are often idle, particularly on night shifts, yet still need to be attentive and ready to take action.

The industrial workplace - and the O&G workplace in particular - is not well documented in the CSCW literature. We have previously suggested that the O&G workplace has parallels to the (commonly-studied) hospital workplace, which may offer opportunity for re-framing existing studies (Heyer and Grønning 2008). The contribution of this paper is an introduction and discussion of the workplace and its work practices, from the perspective of the shift team who work in and across the industrial environment and the control room. In this publication we do not seek to provide 'implications for design', merely to provide a descriptive account of the context and serve as a resource for future design and research.

¹ Names have been made anonymous.

Future work will report on subsequent field studies at other facilities, and develop a set of general design recommendations and insights for the oil and gas context.

The next section of the paper outlines the methodology, followed by an examination of how work is conducted at Kvasir. A discussion section analyzes observations thematically and the concluding section highlights main findings.

Methodology

Ethnographic fieldwork (Hughes, King, Rodden and Andersen, 1994; Randall, Harper and Rouncefeld, 2007) was conducted over a period of ten days, consisting of participant observation and informal interviews. The focus of the observational study was the shift team, however regular day shift employees such as managers and engineers also participated in the study through semi-structured interviews. Observations were conducted by shadowing a particular participant for a shift, or by observing a particular work area, such as the shift leader's office or central control room. Participants would often volunteer information, describing what they were doing and why and speaking aloud information they received on the radio or computer. It was more likely though that we would ask participants questions as they went about their work to provide explanations and clarifications. Eight different shifts across all three shift periods (day, evening and night) were observed, as well as interviews with 28 people. Notes, photographs and sketches were made during observation and interviews, which were reflected upon and expanded at the end of each day. On the following day, observations were discussed with participants in order to validate their correctness or discover alternative or more detailed explanations.

Work Practices

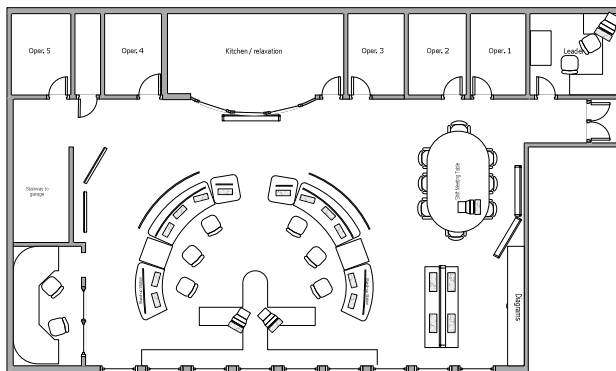


Figure 1. Central control room map. Two control stations are located in the center of the room.

The shift team of around ten people is made up of field operators, two central control room operators and a shift leader and is based in the central control room (Figure 1). Kvasir also has a “day shift” that works regular office hours that includes engineers, administrators and other personnel. The shift leader is a senior member of the shift who manages the team and acts as a mediator between them and the wider organization.

Field operators

Field operators have particular core competencies (such as mechanic or electrician) and are also assigned to particular plant areas. Operators usually work solo as they roam around the plant but frequently liaise with other operators, contractors and engineers. Much of the upgrade and repair work is carried out by contractors, whom the operators oversee, managing work orders and permits as well as ensuring the safety and quality of the work. During a shift, operators also perform maintenance duties, such as draining fluids and cleaning motors.

Field operators have offices adjoining the central control room where they have computing resources and can attend to paperwork. On a usual day shift, operators spend half their time in the plant (Figure 2), half in the administration building. A stairwell from the control room leads straight into a ready room where operators pick up their safety gear and “check in” to the plant with a proximity card. Operators always wear their radio, safety boots and high-visibility, flame-retardant clothes, so they are ready to go to the plant at a moment’s notice. When in the plant, field operators have little in the way of advanced technology except for a radio or frequently, a portable gas detector.



Figure 2. A typical oil and gas plant environment.

Central control room operators

Kvasir has two central control room operator roles (CCROs), each responsible for one of two main process areas. From the control system, CCROs can physically

manipulate the process, such as remotely controlling valves and pumps. CCROs report that there is usually little need for such manipulations, except in the case of major maintenance or during a shutdown or startup procedure. One operator reported performing manipulations in order to better understand the process, “to see what would happen”.

In their role as a hub and mediator of information and control, CCROs exhibit a high degree of multitasking. For example, we observed an operator who was on a personal mobile phone call also juggle multiple landline calls, radio communication with field operators, radio communication with a moored ship and also making process manipulations – all within a two minute period. Each operator sits at a control station (Figure 3) which has a large projected composite display positioned above and behind four individual screens, all of which predominately show process graphics. Process graphics are a visual depiction of some part of the logical process, overlaid with live process values such as pressures, temperatures and flow. The large display’s abstracted, simplified view and physical size makes it easy to see important information at-a-glance. During emergencies or shut-downs the screens are a useful way of keeping a number of people informed about critical parameters without interfering with the CCROs.



Figure 3. A control station.

Smaller displays are set to show areas of current interest and used to perform process manipulations. Surrounding each process graphic screen is an interface to browse or go directly to particular areas of the plant. In some configurations, panels also appear in the display, for example a list of alarms or video camera feeds. Temporary, overlaid dialog boxes can also be shown, most often for detailed trend line views or ‘faceplates’. Faceplates are used to view detailed

information and make adjustments to particular pieces of equipment, with each major class of equipment having its own specialized interface. Three separate computers provide the graphics for the control station's five screens, each with its own keyboard and mouse. This can cause some confusion when moving between the boundaries. Attention can switch quickly to a new screen simply by glancing, while a higher level of cognition is required to release and reacquire input devices. CCROs deal with by attempting to keep the three sets of input devices in spatial relation to their respective displays, to mixed success.

Training, learning and understanding

Field and control room operators both undergo extensive training before being able to work on their own. The training process is a combination of theory learning and apprenticing, with skills being honed and maintained by way of examinations and running through simulated scenarios.

Field operators report that it takes about one year to gain an understanding of a single plant area, perhaps four to five years for the whole plant. Control room operators have a two-year training period before being able to operate a control station solo, and they report that it takes an additional two years to get a good sense of the more complex of the two control areas. Learning periods are dependent on how stable the plant is: an unstable plant is a better teacher than a stable plant. During learning, operators develop an internal model of how the plant works and how the individual components function and fit together to form the process. Both operator roles will have memorized a large number of tags (an identifier assigned to each component of the plant, such as pipes, valves and compressors and used as a uniform referencing system). For a tag number, such as '0300GTFC1AK', an experienced operator knows what it is, where it is located, its history, expected behavior and other properties.

While training provides theory and practical skills, it can still take time for operators to get an understanding, or 'feeling' for the process, a sense of knowing appropriate process values, sounds, smells, vibrations and when things are out of place. Novice CCROs tend to browse process graphics more, as they are not sure which areas to focus on and don't want to miss anything, while expert CCROs tend to jump directly between a small set of graphics. Maintaining an understanding of the process and its operation requires time in the simulator as well as also running the process. For this reason, it is not possible to keep a large number of people 'current' as a CCRO. After returning from the shift's break period (once every six weeks), control room operators run through scenarios in the simulator for around four hours before resuming active duties. Some also read through a log of events to get a sense of what has happened in the plant while they were away. On the first shift back in control, operators report feeling

somewhat overwhelmed by the amount of data and alarms, although using the simulator reduces the period of this sensation.

Inspections

Both operator roles conduct regular observation rounds per shift, using their understanding of the plant's norms to detect irregularities. Control room operators conduct their round by traversing the process graphics with their mouse, looking for anything out of the ordinary. By setting their multiple screens - sometimes up to 13 discrete displays - to show particular areas of the plant, the CCRO can be peripherally aware of activity around the plant. Field operators conduct their round by walking through their part of the process, checking oil levels, looking for leakages and so on. This takes an experienced operator around 15 minutes and is conducted thrice per shift.

The 'pull' mode of fault detection is complimented by a 'push' mechanism. For control room operators, this takes the form of control system alarms (Cauvin, Cordier, Dousson et al. 1998) forcing the operator to examine a particular part of the process for sign of fault. For field operators, the 'push' usually comes via radio, for example a control room operator or contractor making a request. For both, these events occur unexpectedly and can interfere with existing work at hand as they usually require immediate attention.

Ritual and routine

Work is ritualized for safety and quality reasons. As an illustration, consider the scenario when a pump needs to be disconnected so it can be cleaned. One operator (with electrical competence) picks up the isolation form from the shift leader's office and travels to the pump's location at the plant. Once there, she verifies the tag number on the paperwork matches the pump in front of her. She depresses the manual stop button and pins a white copy of the isolation form to the pump. At the appropriate electrical substation, she finds the switch panel for the pump as well as a padlock for the panel. The handle is switched into the off position and the padlock inserted to prevent it being turned on inadvertently. It is critical that the right switch panel is isolated and that it does not get reconnected while someone is working on it. A red duplicate of the form is attached to the panel and the key is put back in the cabinet, or given to the person working on the isolated equipment if they request it. After the inspection is completed, the operator is called down again to reverse the process, step-by-step. Each step is ritualized and formalized so all parties can be sure the process is carried out correctly and safely. The paper copies of the isolation form pinned on the equipment and switch panel serve as visual notifications of the alteration and that protocol is being followed.

Situated action

There is usually some amount of pre-determined structure and activity to operator's work, such as routine inspection or maintenance, depending on the prevailing plant conditions. However, we suggest that most of field and control room operators' work is reactive, based on emergent conditions and activity of others. For example, a field operator might be performing an inspection round, and then be radioed by a contractor to approve that the work site is safe to commence work. The operator interrupts his current task, travels to the contractor and signs one part of the work permit. He then resumes his work, but will be likely interrupted again when the contractor radios again for the final signature when the work is complete. Field operators are also frequently issued commands and requests from the control room, such as to perform isolations, check valves and so on. Because requests are ad-hoc and not centrally triaged, it can result in inefficiencies. For example, we often observed an operator driving down to the plant and start work, but after receiving a radioed request need to travel back to the administration building to pick up a form, and then drive back to the plant.

Collaboration and communication

Kvasir is a small organization of approximately 130 employees with a relaxed, informal atmosphere. The flat organizational structure results in short lines of communication and it is typically easy to directly communicate with the required people, regardless of their organizational division.

Technologies

A patchwork of different technologies are employed for collaboration and communication. Computer-based systems include the ERP system, web-based document management systems, email and instant messaging. Radios are mostly used to communicate between shift members and to on-site third parties such as contractors and ships. Wide-area broadcasting is accomplished using the public address system, installed in the administration building and the plant, however considered ineffective in the plant where loud noise and hearing protection occludes the speakers.

Mobile telephones are pervasive, however because they introduce an explosion risk, are only permitted in the administration building. Some employees carry three mobile phones: one personal, one company-issued work phone and another company-issued emergency phone. Those with an emergency phone are expected to have it with them at all times, even after hours, so they are reachable in case a problem arises. Although most people are reachable through a phone connection (be it fixed or mobile), it is seldom used between people at Kvasir, as face-to-face communication is preferred.

Radios are the predominate medium for mediated communication between shift members. Under normal circumstances, around 20 people are tuned to the main channel used by the shift. In the plant, operators wear helmets with integrated ear protection and headphones, by connecting the helmet to their radio they can reduce background noise and hear the radio clearly. A microphone and button is integrated into the cable so the operator can speak while leaving the radio clipped to her belt. There can be issues with radio congestion, especially during periods of intense activity such as during a shutdown, although operators report that normally it is not a problem. Radios provide a shared awareness for the entire shift, as everyone can hear the broadcast utterances which reflect location, activity and progress of each member. Shift leaders note that a sophisticated presence or location awareness system is not needed because simply listening to radio chatter tells them much of what they need to know.

Cross-talk, static and high background noise can impede hearing of radio communication. Instructions and numbers (such as new set points or tag numbers) are repeated back to ensure correctness. Operators will also physically move away from particularly noisy areas in order to better hear the radio, although normal conversation volume is usually sufficient when speaking. The volume of the radio is continually manually managed, for example in the administration building, operators keep the volume low, selectively turning it up when they hear something that might be important. In the field, a higher volume is used which can surprise operators when returning to the administration building with their headphones disconnected. Likewise, operators sometimes forget they've turned their radio down, with the whisper-volume of the next transmission acting as a prompt to turn it up again. When talking with people face-to-face in the field, operators constantly alter the volume to zone in and out of the radio chatter and balance hearing the person next to them with hearing the radio, another form of boundary management.

Many informants reported not using instant messaging "as much as they should", acknowledging the benefits of instant messaging (asynchronous, quick communication), but preferring a short face-to-face chat. Instant messaging is often used as a support for face-to-face conversation, for example sending hyperlinks and tag numbers through instant messaging prior to visiting someone. It was also used for quick, simple inquiries, such as a control room operator asking the shift leader for a clarification about a work order. Most IM communication took place between employees who had a reasonable amount of personal familiarity. For unknown persons within the wider company, participants reported they were more likely to use email or telephone.

Email was used when a digital artifact was needed to support the communication, such as an attached document, or when a higher degree of formality is required. Email is also preferred when a single communication needs to be distributed to multiple recipients or when an audit trail is desired.

Face-to-face

As a result of a small, informal workplace, a significant amount of communication takes place face-to-face. One informant reports that he prefers to pick up a coffee and meet someone for a chat rather than arrange a formal meeting or communicate via a technology-mediated method. Impromptu “drop-ins” aren’t successful when people are out of the office, however. To deal with this, people check others’ instant message status to see if they appear to be in their office or send a message asking if they have time to chat. Several informants reported being annoyed at not knowing people’s availability well enough: time can be wasted if a person goes to visit a colleague, but she is not there. On the other hand, one informant reported quite enjoying the excuse to get out of his chair and go for a walk, and didn’t see these missed connections as a problem.

Kvasir has an open-door policy, which encourages communication but also disruption. Informants report different experiences with this policy. For some, this open door policy can be more disruptive than it is helpful, as their work is constantly interrupted by a stream of people at their door. Others find it very helpful to be able to directly visit and ask someone a question quickly without having to make formal arrangements. Clearly a balance between these experiences is desired and even those informants who were often interrupted thought it was a worthwhile policy overall. Hallway-facing walls of offices have large glass panes which facilitate awareness of office activity and presence.

Intra-shift

Shift meetings are held at the beginning of each shift at a desk in the control room. The shift leader chairs the meeting and operates the meeting computer, the display of which is also visible on a large screen behind him. Using the ERP system, the shift leader runs through the digital shift log, a list of current issues with the plant. Issues which are new to the shift or have undergone change are focused on, with long-running outstanding issues largely ignored.

Because of the consistent shift composition, shift members recognize each other’s voices on the radio which aids communication. Members know who is responsible for which area for a given shift period, so people are normally summoned by name (“Anders?”) rather than by role or responsibility (“Control room?”). Interaction can also take place wordlessly, in one situation we observed a field operator walk in to the control room and sit next to a CCRO. Without a word, the CCRO, aware that the operator was just working on a pump, switches graphics on a screen to show the pump in question.

Intra-day shift

Formal meetings between day shift workers takes place in dedicated meeting rooms, or rooms with added meeting-support technology. An example of the latter is the maintenance office, which has a large wall-mounted display that the

maintenance leader uses to send a clone of his desktop to so that others can follow along as he navigates through defect notices. Meetings usually take place around computer-based output, be it an on-screen application or slide deck, with the ERP system widely used as a way of navigating and presenting data. There are difficulties with this approach as the program was not designed for presentation, especially not on a large screen.

In a recent change, some Kvasir management personnel now report to management located in distant cities. Collaboration takes place over instant messaging, email, phone calls, and increasingly, video conferencing. Naturally, these types of communication and collaboration are not well suited to the informal style of interaction that is commonplace at Kvasir. Informants noted that it was thus harder to maintain personal ties and feel connected to those located remotely. Travel to either of these locations is also difficult due to Kvasir's remote location and requires a significant amount of time.

Central control room operators

Although the two central control room operators sit in close physical proximity there appears to be little collaboration between them. Partly, this is because when the plant is running as it should, there is little to work together on: they share a role, but rarely share tasks. The alarm list which appears on both operator's large screen (and often duplicated on a small screen too) is shared. As alarms occur, an audible tone is produced and a new entry appears at the top of the alarm list. Operators use the mouse to mark an alarm as acknowledged; collaboration can thus take place around this shared alarm list. Operators do not manage alarms which are unrelated to their part of the plant, however they serve as a useful peripheral awareness as to their colleague's activity and status. Alarms might not be directly related to their part of the plant, but because of the interconnected nature of the process, may have an impact if the situation is not managed. Thus, if an operator notices her colleague's unacknowledged alarm list growing she might ask him how he is going, or simply look over to his desk to gauge his activity. There are different styles of alarm management which can cause some mild tension between operators, for example, some acknowledge alarms as they occur while others prefer to process alarms in batches.

Engineers and the shift

Engineers draw up work orders which are carried out by operators or contractors. A work order might be simple, such as modifying a pipeline's pressure, re-routing flows or more complicated upgrades and repairs. The central coordination artifact for such operations exists in digital form in the ERP system but is regularly printed throughout the workflow. Engineers will occasionally visit the person responsible for carrying out the task with the printed work order, sketches and annotated diagrams in hand, to talk through the plan. Operators have a rich

practical understanding of the process which complements the engineer's theoretical perspective. Engineers noted that operators will often suggest an alternative plan which achieved the same result but was easier to implement. It is important for the engineer to ensure that the operator understands the work plan or shift instruction and some engineers suggested that this was easier to accomplish face-to-face.

Engineers and operators also occasionally work around the stand-up stations, or at the main control stations. For example, an operator might pull up the process graphics for a particular system and use that as a basis for conversation with the engineer, highlighting particular process values using trend lines. The stand-up stations, more so than the control stations, encourage collaboration around the data, and there were a number of occasions observed when operators or operators and engineers worked together around the screens.

Many engineers make a point of regularly visiting the control room to maintain a good relationship with the shift and to pick up on issues unreported through official channels. It also provides an opportunity for the shift, which is largely bound to the control room and plant areas, to informally and directly interact with the engineers. Sometimes these visits might amount to little more than drinking a coffee while hovering near the control station, or perusing process graphics at the stand-up station. More often than not however, conversation would be initiated and news and updates exchanged. One engineer also reported it being useful to talk face to face as some things are not well expressed in written communication. Engineers also note that the frequent contact is useful for building up two-way trust and better understanding of each other's competencies.

Remote communication between operators and engineers is usually via email, partly due work period mismatches. It is also because communication often includes hyperlinks, documents and precise numbers which need to be expressed clearly and in an auditable manner. For urgent issues, operators will send a text message if they don't expect the engineer will check their email soon. Shift leaders also maintain close ties to engineers and will often call them down to the control room for consultation.

Information and systems

In the plant, various informational artifacts are used to situate temporal information to a particular location and piece of equipment. For example, when a valve is manually changed from its normal operating state, a large plastic sign is attached. It has a "blinding" number written on it in temporary marker that cross-references a printed list kept in the shift leader's office as well as in the ERP system. The tag number, date and who made the change is also written on the sign. During maintenance rounds, paper tags are used by operators to keep track of which equipment they have inspected, usually with only their initials or a date

written on them. Once the round is complete (perhaps over several weeks), the tags are taken down. Operational transactions such as isolating an electrical circuit can also result in temporary notices being pinned to different equipment.

The ERP system is the primary common information system (Bannon and Bødker, 1997), integral to most aspects of work at the refinery. For example, if an operator notices a defect at the plant, this is entered as a ‘notification’, which is then reviewed and annotated by the maintenance network, and possibly goes on to form the basis for a work order and work permits. Information is extensively cross-linked, for example when reviewing a work order, it is possible to see if required parts are in stock, retrieve product specification information or diagram the fault in the context of the logical process. New informational artifacts or exchanges are often composed from disparate sources, such as collating diagrams, specifications, notes, instructions and annotations in a work order. Like Fitzpatrick’s (2004) observation of hospital patient records, these composite, multimedia, locally-contextualized artifacts are “living documents” which frame and support work activities rather than being a passive information repository.

Piping and instrumentation diagrams (P&IDs) are perhaps the most important reference document in the oil and gas workplace, or as one operator called them, “our bible”. In the diagrams, a logical view of the process is shown similarly to the on-screen process graphics but in greater detail. They are often referred to when there is a need to isolate a section or to trace through the process. Operators with different competencies read the diagrams differently. For example, some read the diagram for valve details, while others read it for the properties of pipes (such as the type of steel and pressure class).

P&IDs are drafted and maintained on a computer but the canonical version of a P&ID is the printed master copy kept in the control room. They are updated over time to reflect alterations in the plant – these too are living documents. Diagrams are frequently printed, which enables them to be physically attached to a work order, carried to the plant or around the building, or used to support discussions. Reading diagrams online in PDF form is common and considered a useful and viable alternative to paper-based diagrams. Diagrams are usually sparsely laid out and thus do not demand the high resolution that paper affords. They are hyperlinked so that the user can navigate the process or retrieve further detail by clicking hotspots. This is considered one of the more useful aspects of electronic viewing – the ability to explore the process with minimal effort. When using printed diagrams, the user would need to know what part of the plant to print and at what detail, and there is a time penalty of finding and printing or retrieving the diagram if another view is required.

Vignette: Dealing with pressure

To illustrate the collaborative work and use of informational artifacts by the shift, we describe the following event which took place during an evening shift:

Suddenly the irregular audible warnings sounded continuously. Soon afterwards, the pitch becomes higher, indicating the warnings are now alerts. The alarm list on the control station blinks frantically as new alarms are rapidly appended. Across the operator's smaller screens, various displays blink urgently. Two field operators come in from their offices, standing behind the control room operator whose area is being affected. The CCRO had not seen this type of alarm before and was a little worried. He also knew from experience that problems in this process area can shut down a boiler which would be a significant problem. On the control system, he quickly opens trend line displays relating to the boiler, scaling them back in time to look for sudden changes. On seeing that the boiler was not being affected, he relaxes somewhat. One operator returns to his office, the other stays and helps the CCRO as the event unfolds. While the boiler did not appear to be affected, pressure was being built up in the distillation column, which would need releasing urgently. P&IDs are brought up in PDF form on the workstation, and quickly printed out. The operator fetches the printed document and he and the CCRO talk around it, trying to establish what is going wrong and how to reduce the pressure. The CCRO's first thought was to relieve pressure via the flare, but after examining the P&IDs and current process values, it would seem as though several valves are in the wrong positions and that the flow could be routed back into the process normally. They then go to the shift leader, an expert in the distillation area, who agrees with their course of action. A field operator is radioed to make changes to the valves and shortly afterwards, everything returns to normal. The entire activity took no longer than 25 minutes, 10 minutes of which was highly tense whilst the alerts sounded continuously.

As a result of flow re-routing during repair work, parts of the plant that have lain dormant for nearly 10 years were used and an error in the original valve diagram was discovered. After addressing the immediate problem, the CCRO composed an email to the responsible process engineer. He attached a screenshot from the erroneous P&ID with mouse-drawn annotations and included hyperlinks from the ERP and web-based documentation system.

Discussion

Non-digital informational artifacts

As described earlier, temporary in-situ notices are used in order to link status information to the physical plant, for example that a valve has been changed from its normal state, or a circuit has been disconnected. Notices are almost always explicitly linked to a particular piece of equipment by a tag number, the universal referencing system in a plant. Notices serve two main purposes: to make visible that which is invisible and to link the physical with the digital. For example, that a valve is in a changed position is not externally observable without prior knowledge of its proper state. Hanging a sign on it makes it clear to everyone that

it is currently in an altered state. Because of the enhanced visibility, the valve will also be easier to find when it comes time to changing it back. Moreover, notices express *purposeful* action. If a pump was found to be switched off, the presence of a sign hanging on it will indicate that it was intentionally put into that state.

Digital-based information, such as work orders and reported faults are linked to the physical artifacts by way of the signs. Each notice has an identifier in addition to the tag number which allows people to trace why an action was carried out and other particulars. Brief information is also included on the notice itself.

Sticky notes were only occasionally used for inter-shift communication, such as the night-shift leaving a message asking if a pipe should be running at its current pressure. During a shift, paper was used differently depending on plant activity and operator's experience level. Some maintained pads of paper to keep track of process values and reminders, while others remembered everything: for some entire shifts control room operators would not use any informational artifacts beyond the on-screen process graphics.

Most field operators use a pocket notebook of some type - usually rather tattered - in which they keep notes, tag numbers, part numbers, test results, sketches and so on. Notes range from being highly temporal, written down but meaningless after an hour, to notes which were referred back to years later. Notebooks are often used as an intermediary to the control system, an aid for discussion (such as drawing a picture) or memory.

Plant piping and instrumentation diagrams are useful to locate tags and understand the process. A single plant is represented by hundreds or thousands of P&IDs, available digitally but frequently printed out on an on-demand basis. The A4- or A3-sized pages can then be studied and marked up with highlighters and pens, and is often used as the basis of interaction between colleagues. The diagram is eminently portable and does not introduce a safety hazard in the plant where it is used to identify physical assets or as a navigation or memory aid. For example, whilst in the administration building with full documentation available to him, an operator might highlight on a P&ID valves that need altered and then use this annotated diagram as a spatial workflow to carry out the tasks when at the plant. Engineers also use P&IDs extensively. Marking streams, valves, flows and values on the printed diagram is a way of building an understanding of the process and a useful part of the diagnosis process for engineers, whether working alone or with others.

Perhaps the most prominent use of paper is that of work permits. Work permits have a short lifespan – typically no longer than a day – and are linked to a single unit of work being carried out in the plant. They are an important part of the audit trail, to not only ensure that work is carried out properly, but that it is done in a safe and verifiable manner. Usually a contractor carrying out the work will produce the permit, which is first signed by the shift leader. Some types of work (those with a higher risk category) might require additional signatures before

work can commence, such as from the operator responsible for the area. When work is complete, the contractor signs and then the operator responsible for the area signs, taking the permit back to the shift leader's office for archiving.

Signing pen-on-paper is quick and accessible with the only requirement being that the parties and the paper are co-located. Typically, contractors achieve this by radioing for operators when they require a signature, however this can cause delays if the operator is some distance away. Because operators usually need to verify some aspect of the physical scene before signing, the co-location requirement of pen and paper is not necessarily an inconvenience. Co-location can also be achieved somewhat creatively, for example signing forms through open car windows, as one car heads up to the administration building and one heads down to the plant.

Bridging boundaries and flowing activity

Boundaries between systems, people, knowledge, perspectives, practice, groups and locations are bridged in a number of ways. For example, field operators often act as a bridge to the physical process for engineers who do not visit the plant as often. Shift leaders act as bridges between shifts, exchanging information about what has happened during the outgoing shift so the new shift is up-to-date. Artifacts can also act as bridges, in a similar manner to "boundary objects" (Star, 1989; Lee, 2007). Indeed, the sole value of some artifacts appears to be their bridging quality.

Sticky notes are a common example of a bridging artifact in the workplace. Short snippets of information are jotted down and then either pinned up, to inform later shifts, or handed over, for example to exchange a tag number. Their value is often highly temporal, serving a purpose as a bridge in an interaction and then no longer being useful. During maintenance network meetings, work order numbers are often passed to the network leader via instant messaging as the meeting progresses, an even more temporal form of sticky note.

P&IDs serve an important role in bridging the communities of practice (Lave and Wenger, 1991) of engineers and operators, as one engineer reported, "[we] mainly talk with diagrams". The diagram serves as a common foundation upon which new understandings can be built. All parties know how to read a P&ID, and it is through the P&ID that a meaningful discussion can take place, grounded in the common understanding. The diagram also facilitates talking about a process which is physically and logically large and complex through a standardized, simplified proxy.

A common quality of the aforementioned artifacts is their ability to bear free-form expression and direct interaction style. A sticky note and P&ID can have anything written or sketched on top and an instant message can hold free-form text. They are also direct in that they require minimal preparation to use and there

is little interactional work required beyond that to express the desired message. These qualities are not found in Kvasir's web- or ERP-based systems, which typically involve extraneous navigation and form-filling work.

The O&G context brings boundaries to the fore partly because of the rich, multifaceted ecology in which activity takes place. Consider for example, how field operators transition between perusing interlinked information systems in an office environment, discussing a procedure over a printed P&ID through to opening a valve in a noisy, hazardous plant. In addition to organizational and technical boundaries, there are also explicit policy boundaries which partition and restrict access to physical spaces and infrastructure such as electrical and computer network grids. Practitioners flow their activities across the multitude of boundaries and diverse resources in order to accomplish their work. To some degree the work environment (computer-supported or otherwise) supports them in this task, however we suggest that flowing is largely accomplished through assembling and tailoring appropriate resources by workers themselves.

The plant

Shift members have a strong phenomenological connection with the plant, which has also been observed in the nuclear power plant shift members (Vicente and Burns, 1996). They speak of the freedom of walking around in the fresh air, the sound of the plant, the smells: feeling and knowing. Frequent contact with the plant over a long period of time allows them to learn what is expected and what is not, what equipment is prone to failure and so on. Shift members' deep understanding of the physical plant is unmatched in the organization and a useful resource for diagnosis and maintenance.

During inspection rounds, rather than simply examining dials and gauges, operators engage with the plant in a rich, experiential manner. For example, an operator will take his glove off and hold his hand to a motor to feel for heat and vibration. Operators listen to equipment's noise, observe steam quantity, color of flames and smell for gas leakages or burning oil. Because of their experience and knowledge of the plant, operators can in effect sense when something is wrong: for example, there might not be enough steam being produced, a motor might be too hot, or a pump might be rattling. These environmental, ambient cues are an important aspect of the operator's diagnostic and observational role. Manual observations are used to look for faults which are not revealed by instrumentation and thus invisible to control room operators in their control system. For example, it is simple to monitor gradual wax build-up on a valve through visual inspection, however it is only when a valve has seized up that the build-up is apparent via integrated sensors, at which point the issue could be critical.

Field operators visit the plant more frequently than other employees. On a regular day shift, an operator will usually spend about half their time in the plant.

As such they are subject to the variable - and very often harsh - environmental conditions, with frequent rain, high winds and temperatures as low as -15°C . Operators liked the variety of being inside and outside, likening it to having two different jobs. Inside, they enjoy the sociality of control room, outside, they enjoy the freedom of movement and contact with plant. Informants enjoy the physical aspect of the work, being able to hear and smell the plant and “get your hands dirty”. Most control room operators periodically rotate as field operators, and while they enjoy the chance to visit the plant, they dislike the loss of ownership and mastery when they are “only” field operators.

Engineers also note a strong connection with their work and the physical plant. Because engineers do not visit the plant as frequently as for example, field operators, they have a lower level of awareness about their plant area and can easily overlook small-scale faults. As such, they often rely on operators to bring issues to their attention, especially when they visit the plant together. One engineer reports that there is “something special” about being in the plant itself and that doing diagnostics over a video link would not suffice. He notes that a report or account might miss something important and there is not the “personal impression” you get with a firsthand visit. For example, if a leak is reported, he likes to see it to get a sense of its scale, what it looks like and where it is coming from – properties that could be represented in a fault notification, but are unlikely to give the same sense as to actually observe it directly.

Conclusions

The oil and gas workplace presents a rich, multifaceted ecology of action, spaces, information and artifacts in which activity is carried out. While the facility is technologically advanced, various non-digital informational artifacts are used as critical parts of work. In the plant, situated signs are used to express purposeful action, make visible that which is invisible and to serve as a link between the physical and virtual. Pocket notebooks are used extensively to record both fugacious and enduring information. Most safety and quality protocols require paper forms and written signatures, such as work permits.

Boundaries between systems, communities of practice and the physical and digital are bridged a number of ways. Piping and instrumentation diagrams are frequently used as a resource in communication, permitting a higher-level understanding to be reached from the common understanding of the diagram. Pocket notebooks are used to jot down tag numbers or values in the field, which later serve as input when accessing online information at a workstation. Discussions often take place around a particular tag, with sticky notes or instant messages commonly used a way of exchanging these identifiers. Such mundane technologies are well suited to these bridging tasks due to their ability to bear free-form expression simply and directly. As is often the conclusion in CSCW,

designers need to give special attention to the mediation of boundaries in order to support smooth flowing of work across diverse physical and digital resources.

Participants cited a strong phenomenological connection to the physical plant. They note the importance of smell, sound, vibration and sights for understanding the plant as well as detecting and diagnosing faults. More generally speaking, the importance of physical presence emerged as a theme, such as engineers spending time in the control room to build relations with CCROs, or field operators spending time in the plant to learn more about its operation and develop a sense of its norms and quirks. This may suggest limits to tele-operation of facilities, or having operators service multiple plants.

The results reported here hint at some danger in implementing 'obvious' solutions. For example, a system to provide digital P&IDs to the mobile field operator would have limited value over paper P&IDs if it does not support annotations, shared interaction, is explosion-proof and usable whilst wearing gloves outside. In another example, the work permit signing process could potentially be made quicker by using mobile devices which communicate directly with the ERP system and do not require pen, paper or radio communication. This, however, would reduce the shift's shared peripheral awareness of each others' activities since there is no usage of the shared radio channel. Shift leaders and control room operators in particular benefit from the awareness the radio channel provides, even though congestion can be an issue during periods of intense work. An apparatus designed to reduce radio congestion and offer alternative mobile communication to the radio would also need to consider how operators today continually manage volume as they pass between different spaces and selectively place the shared channel at their focus or periphery. For example, if text messaging was to be used to support communication, how would the operator know when a message was important enough to focus on without actually diverting attention and reading the message? How would operators exert fine-grained focus-control?

This paper presented the results of ethnographic fieldwork at a gas refinery focusing on the shift team, which consists of field operators, control room operators and a shift leader. The paper's contribution is an initial description and discussion of the work practices in the oil and gas workplace, a context not well explored in the literature, yet one that poses interesting challenges to the CSCW community.

Acknowledgements

We would like to thank the very helpful and accommodating staff and management at Kvasir for making this study possible. We thank the reviewers for their insightful comments and are particularly indebted for the suggestions regarding flow. The authors wish to acknowledge the support of the Norwegian Research Council and the StatoilHydro TAIL IO project for their continued funding and support for this research.

References

- Bannon, L. and Bødker, S. (1997): 'Constructing common information spaces'. In *Proc. European Conference on Computer-Supported Cooperative Work*, pp. 81-96.
- Bentley, R., Hughes, J. A., Randall, D., Rodden, T., Sawyer, P., Shapiro and Sommerville, I. (1992): 'Ethnographically-informed systems design for air traffic control', In *Proc. Computer Supported Cooperative Work*, pp. 123-129.
- Cauvin, S., Cordier, M., Dousson, C. et al. (1998): 'Monitoring and alarm interpretation in industrial environments', *AI Communications*, vol. 11, no. 3-4, December 1998, pp. 139-173.
- Fitzpatrick, G. (2004): 'Integrated care and the working record', *Health Informatics Journal*, vol. 10, no.4, pp. 291-302.
- Heath, C. and Luff, P. (1992): 'Collaboration and Control - Crisis Management and Multimedia Technology in London Underground Line Control Rooms', *Computer Supported Cooperative Work*, vol. 1, pp. 69-94.
- Heyer, C. and Grønning, I. (2008): 'Cross-workplace perspectives: relating studies from hospitals to an oil and gas workplace', in *Proc. NordiCHI'08*, pp. 475-478.
- Hughes, J., King, V., Rodden, T., and Andersen, H. (1994) 'Moving Out from the Control Room: Ethnography in System Design', in *Proc. Computer Supported Cooperative Work*, pp. 429-439.
- Lave, J. and Wenger, E. (1991): *Situated Learning: Legitimate peripheral participation*, Cambridge University Press, Cambridge, UK.
- Lee, P. C. (2007): 'Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work', *Computer Supported Cooperative Work*, vol. 16, no. 3, June 2007, pp. 307-339.
- Martin, D., Bowers, J. and Wastell, D. G. (1997): 'The Interactional Affordances of Technology: An Ethnography of Human-Computer Interaction in an Ambulance Control Centre', in *Proc. of HCI on People and Computers*, pp. 263-181.
- Randall, D., Harper, R., and Rouncefield, M. (2007): *Fieldwork for Design: Theory and Practice*, Springer-Verlag, New York, USA.
- Star, S. L. (1989): 'The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving', in M. Huhns, (ed.): *Distributed Artificial intelligence (Vol. 2)*, Morgan Kaufmann, San Francisco, USA, pp. 37-54.
- Vicente, K. J. and Burns, C. M. (1996): 'Evidence for direct perception from cognition in the wild', *Ecological Psychology*, vol. 8, pp. 269-280.

Asking for the moon *Or model-based coordination in distributed design*

Kjeld Schmidt, Hilda Tellioglu*, and Ina Wagner*
Copenhagen Business School, *Vienna University of Technology
*schmidt@cscw.dk, *{ina.wagner, hilda.tellioglu}@tuwien.ac.at*

Abstract. This paper reports on a study of practitioners in engineering design striving to transform their work practices so as to be able to cope with complex interdependencies across global production networks. As a key feature of these budding coordinative practices, practitioners are trying to build computational ‘models’ of the ‘design space’ of their enterprise. The paper examines the difficulties they face in developing these models.

Introduction

Ongoing changes in the global political economy seem to be accompanied by concomitant changes in the organization of cooperative work in enterprises and institutions. This transformation is, perhaps, particularly pronounced in manufacturing.

For most of the 20th century, the activities of engineering design and production in manufacturing were typically organized within the framework of vertical corporations controlling more or less the entire process from extraction of materials to final product assembly and from design to production (Chandler, 1977). By contrast, the process is now — increasingly — ‘fragmented’, to use the expression adopted by economists studying the phenomenon (Arndt and Kierzkowski, 2001; Cheng and Kierzkowski, 2001). The pin-making process described by Adam Smith (1776), in which the craft work of manufacturing pins had been decomposed into a dozen of specialized activities each of which were allocated to a particular workman, has, so to speak, been disassembled and dispersed over a range of specialized enterprises in different locations. Consequently, a large and steadily

increasing part of (national and international) trade consists of trade in (simple or composite) components as opposed to final products. That is, on one hand the entire manufacturing process is now being distributed over multiple — sometimes thousands — of enterprises. On the other hand, the constitutive units become increasingly specialized. What emerge, then, are global production networks (Arndt and Kierzkowski, 2001; UNCTAD, 2002; Berger, 2005; UNCTAD, 2005). The topologies may vary; some may look like ‘supply chains’, others like hierarchies of thousands of small enterprises controlled by a transnational corporation, and others again like proper networks.

A variety of motives are of course at play in this transformation process. In many cases the driving motive is that of reducing the cost of labor by outsourcing to countries with substandard labor conditions. However, other motives, less transient and more sustainable, are also involved, such as the advantages of increased specialization, economy of scale, etc., made possible by the radically reduced costs of transportation and communication (cf., e.g., Harris, 2001; Levinson, 2006).

Whatever the motive, the ‘fragmentation’ of the design and production process — *i.e.*, its increasingly distributed character — raises acute coordination problems for the participating cooperative ensembles.

The reasons for this are rooted in the nature of design work. In his classic analysis of design work from 1964, Christopher Alexander argues that ‘What does make design a problem in real world cases is that we are trying to make a diagram for forces whose field we do not understand’ (Alexander, 1964). That is, design is a ‘wicked problem’, to use the term suggested by Rittel and Webber a few years later: ‘In order to describe a wicked problem in sufficient detail, one has to develop an exhaustive inventory of all conceivable solutions ahead of time. The reason is that every question asking for additional information depends upon the understanding of the problem — and its resolution — at that time. Problem understanding and problem resolution are concomitant to each other [... The] process of solving the problem is identical with the process of understanding its nature’ (Rittel and Webber, 1973).

Christopher Alexander went on to claim that ‘more and more design problems are reaching insoluble levels of complexity’ (Alexander, 1964). If this was not obvious when he wrote it, it is evident now, as networks of industrial enterprises struggle to master distributed product design: ‘These complexities are compounded drastically when solving a “wicked problem” involves multiple actors, in that different aspects of the problem are addressed by different designers and the interdependencies among these aspects, and hence between the actors, emerge and change as the design project unfolds’ (Schmidt, 1998). When design work becomes distributed over global networks of specialized enterprises, the problem becomes malicious.

Our study focuses on describing this challenge to engineering practitioners in production networks and how they try to cope with it.

The study

We ground our arguments in extensive fieldwork carried out in two companies in the automotive industry: *Newcars*, an automobile manufacturer, and in particular *Carparts* on the supplier side. We engaged with these two companies as part of EU Project MAPPER whose objective it was to develop, introduce, and evaluate an approach to ‘model-based adaptive product and process engineering’.

In this paper we focus on *Carparts*, which belongs to the 2nd tier suppliers of the Automotive Supply Chain. It produces ‘seating systems’ (climate control, motion controls, etc.) as well as head restraints, control cables, and gear shifts. It faces problems in managing myriads of highly interdependent tasks in a distributed network of suppliers. It also seeks to improve its ‘process of innovation’, with a view to developing and evaluating design alternatives for its products.

Empirical material was collected at *Carparts* during two field visits, each lasting several days, in November 2005 and March 2006, with the purpose of trying to ensure that technical requirements be grounded in actual work practices and needs at the user site. During these field visits we had the opportunity, through ethnographic methods, to study a series of activities related to advanced engineering in the company. During our first visit we were able to observe how projects are managed. We followed co-located and distributed meetings, project meetings as well as design reviews, and ongoing work at a series of workplaces in design, testing, and purchase. During our second visit we focused on practitioners’ interactions with external suppliers and on the company’s ways of managing projects set up specifically for product and process innovation. In addition to this fieldwork, we engaged with various staff in a series of interventions. One of the authors also participated in a number of modeling sessions carried out by the MAPPER modelers with project responsables from *Carparts*. Our final involvement with *Carparts* was a validation event in November 2007, where we, among other things, were exposed to an approach to product modeling which the internal project manager for MAPPER had developed. On that occasion, we also observed a modeling session dedicated to the creation of a model of collaboration with suppliers including a demo of model-support of customer-supplier design collaboration.

The long-term collaboration with personnel at *Carparts* (which was further strengthened in project meetings of all sorts) allowed us to acquire substantial knowledge of the ways of working in this company and its problems. But we also need to emphasize the limitations of our fieldwork with regard to the use of the modeling approach promoted by MAPPER. The models we will describe are constructed as part of experimentations and have not yet been deployed. They were

developed over the course of almost one year by practitioners (in different professional and organizational roles) in collaboration with modelers (consultants, researchers, as well as in-house specialists in modeling). We were not able to actually observe the day-to-day process of modeling but rely on presentations of this process by the internal project manager. However, we have witnessed some the difficulties of those involved in producing these models and the numerous conflicts surrounding this process. Hence, when we refer to ‘modeling’ in this paper, we do not describe an already existing practice. What we look at (and document) are practitioners’ attempts to develop such a practice and the associated techniques as well as their problems with doing so.

A view from the top

The automotive industry launched on a large engineering outsourcing activity in the late 1990’s. This had strong implications for the integration and coordination of knowledge and competencies on the one hand, for the organization of product development on the other hand. Companies such as Toyota, Renault or Fiat implemented the concept of ‘platform’ – units that are based on a core team formed by several professional profiles who follow the whole life cycle of a product. Hand in hand with this a modular product architecture was introduced (Bonazzi and Antonelli, 2003). The strategy was, and still is, to separate component design from developing the concept for a new vehicle. Component design starts well before the concept design for a particular car and involves a panel of what is called 1st tier suppliers. These are strategic partners who are actually involved in co-design and substantially contribute to product innovation (Midler, 1995). Another category of suppliers are those of parts with an influence on styling and where also a high innovation rate is expected – lights, seats, windows, electronics, hydro-forming, etc. This engineering outsourcing activity has been described as a move towards the car company becoming a ‘systems integrator’ (Becker and Zirpoli, 2002).

Our case study at *Newcars* focused on one central phase of cooperation with suppliers, the so-called ‘target setting process’. At the beginning of the development of a new car there is the ‘vehicle concept’. As part of this, desired product properties or ‘targets’ are formulated on the basis of a market analysis, interviews with customers and/or a focus group, an evaluation of the competitors, etc. The aim is to identify the main features of the product in terms of security, comfort, sportive performance, price range, climatic comfort, etc. This is also called the ‘voice of the customer’. A ‘performance tree document’ is created which lists the features starting with top-level requirements. Target setting is led by the marketing people in collaboration with engineers. Qualitative criteria for each feature have to be translated into technical criteria and parameters, e.g., system efficacy, or air distribution. Also the price has to be set for each of these features. Different

types of engineers are involved: ‘performance engineers’, most of them with a background in Computer-Aided Engineering and virtual testing, have to set targets and perform the first analyses, in collaboration with engineers responsible for systems of physical components (‘RdS engineers’), who have to decide whether these targets are feasible. The negotiation of performance is a complex process involving a large number of suppliers with whom targets are discussed and if necessary modified. The aim is to have modifications fixed in the early phase, since the cost of engineering changes increases as development advances. This is a process of optimizing performance and integration over all vehicle systems and parts. There are often conflicting targets and always conflicts with cost targets.

This process is supported by a range of IT systems. The PLM system (Product Lifecycle Management) contains pertinent information concerning product development (engineering Bill of Materials, CAD drawings, a digital mock-up environment for virtual testing, a specification of the formal process of engineering changes, and so forth), but it only supports the engineering aspects of target setting and not the requisite communication and coordination with the supplier. In fact, within the *Newcars* Group different PLM systems are in use. Similarly, the system for managing the performance tree is proprietary and thus not shared with suppliers, only the SSTS or Sub System Technical Specification system is. Not surprisingly, updating the State of Requirements document after each target modification process takes time. Hence, while formal communication is mediated by the SSTS, day-to-day interactions with suppliers are done by email, phone and, if this is possible, through shared CAD documents. The complexity of this process together with the high dependency with suppliers, many of whom are chosen by Purchasing, creates huge management problems for *Newcars*.

A view from the middle ...

From the point of view of a particular work organization in the middle of the network, or rather enmeshed in the middle of multiple networks, the whole thing is even more complex. On the one hand it produces components for a range of customers, often-large corporations, and on the other hand it is itself a customer of a network of suppliers. When components are highly standardized items, commodities, this position is classic and does not pose a particular challenge. Nor do very stable ‘supply chains’ pose a major challenge to participants. The challenge arises and becomes a major one when component designs are not standardized and stable; that is, when customers request different and varying design configurations. The enterprise-in-the-net is then exposed to conflicting force fields. From its customers it is presented with requests and requirements with respect to its products that it will have to find economically viable solutions to: ‘Can we do this at all?’, ‘Do we have a design we can modify?’, ‘Do we have to open a new product line and could we then reuse the new design for other purposes in some modified

form?', and in any event: 'What will it cost?' and 'Can we meet the schedule?'. And conversely, as far as its own suppliers are concerned, the enterprise-in-the-net of course poses the same requests and requirements. (New design options may of course flow in the opposite direction, 'up stream', just as legacy design options may disappear from the pallet, for instance for reasons of environmental protection).

Coordinating with these different stakeholders is difficult. It involves, for example, negotiating specifications with several suppliers while routinely resolving the problem of aligning different part-code naming standards, and so forth. Moreover, standards in manufacturing differ across national boundaries and we have witnessed several meetings at which such mundane differences created severe problems.

U36201 Project meeting no: 12		Agenda:
Date: 2005-11-03		1. General issues
Participants: R&D, OGR, TOEK, JOL,		2. Sales
POB, PEHC, ANNV, MNJ		3. R&D, testing
Distributed to: MRAH, FRED, MAD, L&U,		4. Purchase
GGC		5. Production
		6. Timing

General issues		S-release: w.541	PPAP: w.645	SOP: w.717
Work no.	Issue			Assign.
38.3	Written confirmation of plastic covered brace is ok, appearance approval. Open issues from design is closed with JCI.			DIASV
43.1	Trimming prototypes ordered from JCI. To be supplied w.544			DIASV
Info	U-36200 Bar-Carrier project will be implemented in the ER project.			

Sales				
Work no.	Issue			Assign.
39.2	Does JCI accept our soft tool quote? Customer accepted specified soft tools in design freeze meeting.			RJO
42.1	Quote price and tooling price impact on collar holes instead of welding nuts. Await furnish ER from JCI.			RJO
43.1	Push for ER due to styling change.			NO/DASF
43.2	Assembly of EFP by JCI on outer HEC? Collect arguments to give a negative answer.			RJO

Figure 1: Issue list.

At present, managing this complexity relies heavily on documents that have been pre-defined for each project stage and that are meant to ensure 'best practices' as well as accountability. For each stage in the project, the project needs to pass a 'gate', at which point the project manager is supposed to have the required documents ready. This is checked manually by the Steering Committee coordinator and there is a formal signing-off of each 'stage gate'. The standard format for documenting technical information is Excel files. In these documents information is arranged in the form of lists of parts, materials, or tasks organized according to different principles. These lists are produced and used by engineers and their project managers.

A key document in the hands of a project manager is the so-called 'issue list', which is central to handling the weekly project meetings (Figure 1). Each issue

list has a header with the project name, meeting number, date, the list of participants, the list of people to whom the list is to be distributed, and the agenda. The form of an issue list ensures that issues are addressed in a particular way. For each issue the list specifies activities, responsible persons, and deadlines. Issues are identified by the number of week in which they have been addressed and a short text. Starting with general issues, most lists we encountered represent, in a rather loose way, a certain order of priority and/or different actors (e.g., R&D, purchase, sales) and/or project stages (e.g., quoting, testing, releasing). There is a particular meeting dynamics around issue lists. At the beginning of the meeting the project manager opens the issue list. S/he addresses each issue, step by step, asking for status information, changing parts of the task description or the deadline when relevant. S/he may also introduce additional issues, specify actions, and so on.

Issue lists are at present the main means for evoking and advancing open issues in a project at *Carparts*. It is also the main means for dealing with uncertainties in as much as the issue list allows practitioners to project complex and difficult issues onto separate and linear tasks, expressed in terms of concrete and simple steps. The list also ensures accountability in that commitments are specified and can be traced as it is made transparent which week a decision on which issue was taken. We can say that the main function of the issue list is to document issues and the related decisions for purposes of awareness, reference, control, and accountability (Jacucci, *et al.*, 2007).

However, there are numerous problems with this ‘document-driven’ way of managing work. Since there is a host of documents ‘behind’ the issue list that needs to be aligned, updated, and shared within the network, much cross-checking, for example, has to happen in the process of negotiating specification parameters with multiple suppliers. To put it bluntly, as it is now the material specification process is unbelievably cumbersome and tedious.

For example, Jill is working on the specifications for a heating wire, a new product. She has improved the specifications step-by-step, consulting with the supplier. She now finishes the third release of the specifications for wires of different width to send it off for signatures to Design and Production. To register a new issue she has to pick an issue number from one of the folders located in the main building. This is a serial number that is totally unrelated to the part number or specific task. Jill signs and enters the date. At the moment she, in consultation with the supplier, specifies the ‘bare single diameter’ because this is a piece of information that the design department needs. For this purpose she examines repeatedly an email sent by the supplier who has specified the nominal weights of enameled products for her. She also changes various text strings such as ‘bunched and reinforced’. At various points she brings out her calculator, checking a value. Jill has to go through each single line in the five documents describing five wires with different width (and part numbers), checking carefully. She then creates a PDF file, inserts ‘sign this document’ and crosses out the part number on a small

hand-written list. It takes her almost five minutes to attach all the documents to be sent off: she opens each document to see if it is the right one, even though the file names indicate the part and issue numbers.

To better deal with processes such as the one we have described, *Carparts* has initiated the introduction of a document management tool (PLM), but the introduction is already delayed and has resulted in much frustration with what personnel at *Carparts* perceive as a pressure to produce more and more documents ‘for others’. In parallel practitioners started experimenting with modeling as a way of capturing complex interdependencies and, eventually, making processes, such as material specification work, more efficient.

Modeling the design space at *Carparts*

An enterprise-in-the-net such as *Carparts* may, over time, wind up in a quagmire of proliferating product models and variants that will completely neutralize the benefits of specialization and economy of scale. To counter that, such enterprises need to ‘map out’ the design space, that is, the extant product portfolio (models, variants, alternative components and materials), the design parameters for each product model (i.e., that which can be changed), and the interdependencies of the different design parameters, e.g., ‘If you do this, then you also have to change that’.

This mapping effort is a daunting task. It is a cooperative effort of significant complexity, as it involves engineers, designers, production managers, marketing people, etc., who obviously represent different professions, different conceptual worlds, different economic and organizational interests, etc. This would in itself make such cooperative mapping effort of interest to CSCW. But not only that: it is an effort that in the eyes of practitioners themselves might benefit greatly by computer support based on computational representations of interdependencies of design decisions and design tasks, that is, computer support of a kind that is central to CSCW’s concerns. This issue was on the agenda of CSCW from the very start and has been pursued under labels such as ‘common information spaces’ and ‘organizational memory’. For good reasons much of this research has focused on the domain of technical design (cf., e.g. Conklin and Begeman, 1988; Subrahmanian, *et al.*, 1994; Subrahmanian, *et al.*, 1997).

What we have observed, however, is that practitioners have, in a strictly experimental manner, actually begun building computational design space maps, or ‘models’ as they term it.

Now, *modeling* is a concept that is fundamental to engineering competencies but that is apt to mystify the uninitiated, as vividly described by Pepper White in his account of his miserable student years at MIT: When a teacher explains that ‘before you can control a system, you need to control the performance of a system’, but ‘once you know how to model things, you can model anything’, White,

perplexed, thinks to himself, ‘Model. Model. Model. Eventually I’ll be able to use that word without blushing’ (White, 1991, p. 121 f.). Ultimately, however, White begins to understand the concept: ‘Model. Key word. So an abstraction is like a model. And a model of a system may be composed of linked models of smaller systems, or subsystems’ (p. 218). — No surprise then that engineers, faced with the challenge of configurable design in production networks, would approach the problem as one in need of ‘modeling’.

While a ‘key word’ in engineering culture, the term ‘model’ is a source of ambiguity in that different stakeholders use the words ‘model’ and ‘modeling’ differently. Models of different kinds in fact abound in the industrial world, typically engineering models (energy flow models, mass transfer models, etc.), but also process and product models and models of organizational structure, workflows, etc. Professional ‘enterprise modelers’ on the other hand talk about ‘powerful’ models in support of collaborative business networking. The models we talk about here and that practitioners at *Carparts* are now asking for, are not engineering models and so on but models of the *design space*, that is, computational models that may reduce the cost and effort required to manage the design space, including design options, costs, process of innovation, etc. as well as the concomitant *workflow models*.

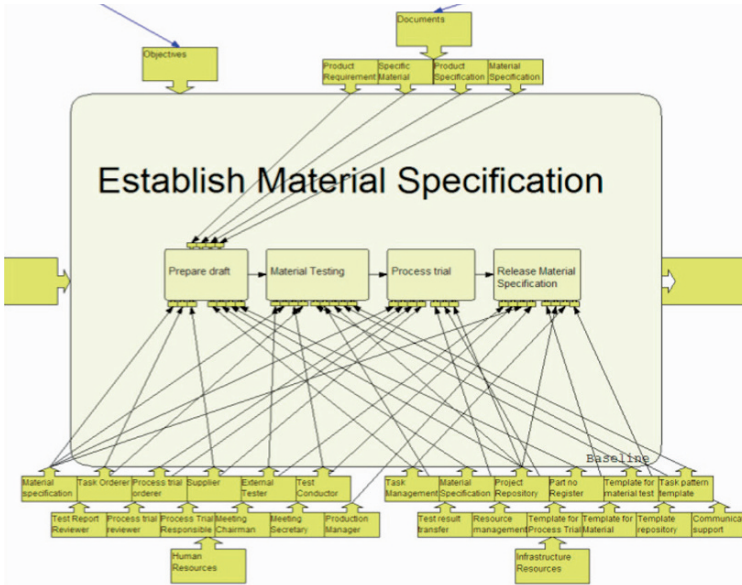


Figure 2: An example of a task pattern: material specification (fragment).

The ‘modeling approach’ was introduced at *Carparts* in a situation of increased pressure for ‘innovation’ (several projects had been set up to ‘improve the process

of innovation’) and it competed from the start with the not yet implemented document management system PLM.

An approach to ‘enterprise modeling’ named Active Knowledge Modeling (AKM) was presented to the project team at *Carparts* - ‘knowledge sharing’ computational platform and associated ‘methodologies’ designed for the purpose of mapping relationships between products, organizations, processes, and systems of an enterprise (Lillehagen and Krogstie, 2002). The team at *Carparts* created several models using AKM during the first two years of the project. These resulted in so-called ‘task patterns’ for, e.g., the material specification process, which we briefly described (Figure 2), but also for more generic activities such as preparing and conducting a meeting. However, working with these task patterns in a small pilot trial did not convince practitioners. Not only was the tool difficult to handle due to a not well-designed user interface; but working through highly detailed sequences of tasks proved cumbersome. After a long debate it was decided to focus on product modeling and at the same time to provide a new interface called Configurable Virtual Workplace (CVW). The main idea was to connect product with process descriptions and to support practitioners in arranging tasks and subtasks connected to their own specific activities, such as design or production, around the product-in-attention.

Modeling a product requires what in modeling jargon is called ‘externalization of product knowledge’. This can be done on different levels: by expressing concept and solution principles, properties and parameter structures, functions and services, systems and capabilities, forms and features, material and appearance, location and spatial relations, environmental aspects, costs and economic concerns, legislation and standards, production and maintenance, life-cycle and end-of-life considerations (Carstensen, *et al.*, 2008). A ‘complete’ product description, or so runs the argument, facilitates working with ‘views’ that focus on the aspects needed for the current work, while ensuring consistency across views in a comprehensive manner.

While the general ideas behind this approach seemed clear, it took practitioners at *Carparts* some time to ‘discover’ how to build useful product models and what to do with them. The experiences we describe are the outcome of a process that was driven by the ‘use-case manager’. We call him Paul. Finding the initial modeling sessions within MAPPER unsatisfactory, he was delighted when he came across a PhD thesis on product modeling for configurability in manufacturing: he scrutinized every page and began producing small conceptual models, first with Excel, later with the MAPPER modeling tool. He set up a small user group, including a CAD technician and two interns, and they began working, undertaking on average one modeling session per month. The idea was to create a complex product description by decomposing the product into Configurable Components (CC) as well as material ‘requirement components’, and to attach to each of these

components a set of validated variant parameters, product properties range, and interface requirements (Tellioglu, 2009).

The team decided to start with simple examples and to work their way bottom-up to more complex product descriptions. They chose to work with seat heating and first spelled out the seat heating conditions and alternatives for the requirement ‘avoid cooling’ (Figure 3), systematically listing all relevant parameters. Paul describes how difficult it was to agree on the parameters that define product variants: ‘We have been spending a lot of time [trying] to identify what in the product variation should be modeled as a performance parameter [PP], what should be modeled as a design parameter [DP], what should be modeled as a constraint parameter [CP], and what should be modeled as a variant parameter [VP]. And there were no real guidelines of what is what’.

Seat heating conditions and alternatives			
		VP1: Climate	
		VP2: Average travel time	
FR1: Avoid cooling	DS1: Heat surface of the seat		DS2: Insulate the body from the seat
PP1: Price	CP1: Flammability		
PP2: Weight	CP2: Fogging		
PP3: Heat losses			
PP4: Comfort	DP11: Price of materials	DP21: Price of materials	
	DP12: Density of materials	DP22: Density of materials	
	DP13: Heat power	DP23: Insulating effect	
	DP14: Heat distribution	DP24: Moisture diffusion	

Figure 3: Product description in spreadsheet document created by Paul at *Carparts* (fragment).

The next step was to create a model of ‘seat heating’. When designing a seat heater, materials have to be chosen for the carrier, the assembly glue, and the seat heating conductor (see Figure 4). As regards the carrier material, the team identified elasticity, environmental footprint, and cost as the main factors, and polyester fabric plus foam and polyester felt as the currently available materials. Having gone systematically through these requirements and confronting them with the currently available choices, their conclusion was that there was a missing combination of properties on the market – ‘PUR free and highly stretchable’ - and that *Carparts* might have found a carrier material they could sell, since none of its competitors uses it.

Paul’s story goes on with the modeling of other components, such as the glue material and the heating wire. For the latter, the heating wire, requirements or properties (electrical, mechanical, failure modes, cost), design parameters (conductor material, thickness of strands, surface layer, fiber reinforcement) were defined, and the option of serial versus parallel circuits was identified as an additional parameter influencing the choice. In this case the choice was between sinus

wires and alloyed wires. Here their conclusion was that ‘very thin copper strains with fiber reinforcement would be ideal’. However, constraints of production have to be taken into account: ‘... *but in reality we are using the alloyed wires, because [of] the constraints in production: the wire layout with the fiber enforcement wire is not doable. And if you add in PVC and insulated wire then there are constraints in peeling insulated wires; so then you will damage the wire*’ (Paul). He took this as evidence for the fact that design choices are interdependent and may have repercussion for production: ‘*So this [is] why we say that the configuring of [a] product should be extended also to configuring the production*’ (Paul).

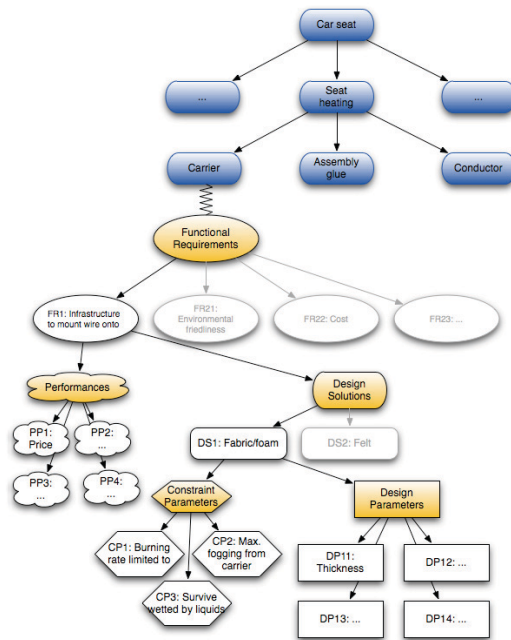


Figure 4: Model of the product part ‘carrier’ at *Carparts* by using Configurable Components (CC)

These more conceptual exercises motivated practitioners to take a step further. The simple product model was enriched with a small executable part that should help them probe how to support their collaboration with suppliers (Figure 5). The scenario was the following: Jane, the responsible for material specifications at *Carparts*, opens the graphical workspace of the modeling tool, searches for a specific wire and enters a specification for resistivity. As a result, the customer responsible at the supplier involved is notified; he edits his own specifications in his own Excel file, and Jane then receives and reviews the result. In this way a specification document is built, turn by turn, both parties seeing exactly the same data while the ‘secrets’ of each party (ownership of data and formulas) are safe-

guarded. In this scenario the responsible for a specific family of materials can also see the aggregated values and compare them to the customer request.

Name	Unit	Min.	Max.	Nom.	Tol.(%)	Status
Conductor cross section	mm ²	0,172	0,262	0,171		
Length / weight	m/kg					
Number of single wires		42	126	84		
Outside diameter	mm	0,025	0,083	0,054		
Single wire diameter	mm	0,05	0,05	0,05		
Tensile strength at RT	N	10				
Wire resistivity	ohm/m	0,963	0,233	0,146	5	

Figure 5: Editing specifications in distributed teams at *Carparts*.

Other practitioners saw this simple executable model as a good checklist: ‘you can see the status [of the specifications] and can highlight the risks from the beginning’. They were not only interested in adding more details to the model but pointed to the value of implementing design rules – ‘there are so many rules around wires – so as to be able to replace the [spreadsheet] tools we are working with now’. They could also imagine using the model for prioritizing sales options: ‘Sales is very impatient, even for early quotations; if they have the tool they could see for themselves’.

The importance of managing complex design decisions within *Carparts* in relation to its numerous suppliers was also highlighted by a quality problem that had come up the year before: it seemed a ‘hot topic’ at that time but got forgotten when the responsible employee left the company. The problem had to do with the quality of the lamination between two foam layers and had been noticed by one of their customers. Thorough analysis of this case had made it apparent that there had been a failure on the side of *Carparts* in communicating certain material specifications to one of its suppliers, which had several repercussions. Paul used this example to propose a model that captures the status of requirement specifications with different suppliers.

Paul comments: ‘This is an important lesson for the product modeling [effort]. If we don’t catch the product requirement, then we will not catch the business agreement either. Once it comes to things, once it appears false in the deliveries, then you must know whom to blame. And these specifications are the basis for

[deciding] whom to blame. So I mean, if we have a better specification of the products, then we are safe in our business agreement.’ His idea was that ‘alarms’ for missing requirements could be built into the model (e.g. ‘no action needed’ (green), ‘start negotiate the requirements’ (yellow); ‘request missing requirements’ (orange); ‘start develop new solution or don’t quote’ (red)).

Paul’s story describes a progression from, at first simple, ‘conceptual’ product models to, still also quite simple, executable models based on these product models. It also indicates a diversity of open questions. A key question is which properties of a product to make visible. As the small examples show this is dictated by practical concerns. What the relevant design parameters of the heating wire are – ‘conductor material, thickness of strands, surface layer, fiber reinforcement’ – results from the current practice at *Carparts*. But it also depends on where *Carparts* thinks they can innovate, or where they think one of their suppliers could contribute something, or on new requirements, such as the EU directive concerning lead-free components. When building such models the critical issue is to capture the relevant permutation options: to which extent can practitioners rely on the completeness of these product descriptions in a model that has been constructed for specific design purposes?

CSCW research has quite early pointed out that ‘the cooperative ensemble reproduces the multiplicity of its environment in the form of the multiplicity of “small worlds” of professions and specialties’ (Schmidt, 1991, p. 6). Hence the challenge of bringing multiple, incommensurate perspectives together. As Paul described, it is this incommensurability that is so hard to resolve: ‘where he [a modeler] used variant parameters, he should have been using performance parameters’, and so forth. For example, addressing the question of how many configurable components to define and on which grounds to decide this, Paul observed that a supplier has other ideas about what to maintain as configurable than has a car manufacturer for whom it is the car part as a whole that is of interest. All these decisions are by no means arbitrary, but they become exceedingly difficult when multiple perspectives are involved.

The issue, we find, is that selecting parameters depends on the particular perspective that practitioners apply and the context for which it is needed: ‘there is no best model’, somebody remarked. In the sciences and in engineering, modeling is a (typically quite systematic, sometimes rigorous) procedure of abstraction for creating useful representations of aspects or sections of the world. It is purposive, therefore internal to a specific practice. No model of a given section of the world is ‘true’ or ‘false’ in splendid isolation from the practices to which it belongs. Rather, models are ‘useful for the purpose’, or ‘not so useful’, as the case may be. — ‘Useful for the purpose’, but for which purpose? Different practices (e.g., concerning production and procurement of insulation, wiring, adhesives, fibers, as well as sales) are characterized by different concerns; they address different aspects or sections of the world with different structural and dynamical characteris-

tics, and they thus conceive of the world differently, apply different criteria of importance, success and failure, etc. Consequently, when it comes to modeling, practitioners of different branches of engineering design have different perspectives that in turn indicate a notion of central object or ‘unit of analysis’ as well as criteria of what to ‘foreground’ and ‘background’ in modeling.

Moreover, even within a given perspective, relevant trade-offs dictate preferences in modeling commitments. The top level trade-off is ‘what’ in the entire world to include or not include in the model’s explicit representation, depending on the costs of handling (gathering, eliciting, validating, maintaining) the requisite information in the model, versus the advantages gained by using the model. Other crucial trade-offs exist in structuring the model, especially in the choice of level of ‘granularity’ (level of detail) and of ‘specialization’ (depth first) versus ‘multiplicity’ (breadth first).

These concerns run deep and cannot be dealt with once and for all. They are here to stay. However, there are also severe limitations with current modeling notations and techniques that may, conceivably, be resolved or amended.

Existing modeling notations are quite generic. There are first of all difficulties with expressing modeling primitives and relationships at the appropriate semantic level, that is, in categories such as, e.g., ‘part/whole’, ‘cause/effect’, ‘pending/decided’, etc., as opposed to the highly abstract categories of the object-oriented paradigm such as ‘object’, ‘class/member’, etc. To overcome these limitations, modelers have introduced the notion of ‘templates’. As opposed to the generic notation of object-oriented modeling, ‘templates’ offer a specialized notation and a library of specialized objects and relationships that have been predefined in a ‘meta model’: ‘the specification work can be significantly reduced by describing the manufacturing or logistics system by a re-usable template, and store it within a library for later use [...Structuring] the templates in an object oriented class structure saves modeling effort and at the same time supports additional transparency as well as some standardization’ (Rabe and Jaekel, 2002). Behind the ‘templates’ are different ‘approaches’, such as POP* (Process, Organization, Process and System), ICOM (Input, Constraint, Output, Mechanism) and CPPD (Collaborative Product and Process Design). (For an informative review of ‘process modeling languages’, cf. (Mili, *et al.*, 2003)).

The choice of template obviously determines what kinds of relationships (hence perspectives) can be modeled (hence expressed). For example, during another modeling session in the project a modeler explained: ‘*Part of planning and setting up a modeling environment is to select the right kind of modeling template, the right kind of modeling languages. But [most] likely, since you can add new things later, depending on the needs as they arise, it is rather flexible as well. You can start modeling using simple templates and add as things go along*’. Choices were formulated in terms of template names, such as ‘in this case I think we will

use ITM [Information Technology Management] or BPM [Business Process Modeling]’ or ‘so we should use a BPM template and not a CPPD template’.

More debilitating, however, existing modeling approaches implicitly presume hierarchical topologies and thereby seduce users to artificially try to enforce orthogonal distinctions onto other forms of relationships. This makes it exceedingly difficult to express complex interdependencies. This limitation may be related to the presumption that relationships necessarily must be represented in the form of two-dimensional graphs in order to be ‘user-friendly’. This assumption may turn out to be a prejudice.

Finally, given the enormity of the challenge of building computational design space models, whence the rush? We have no way of answering the question. But some explanations seem likely. Firstly and obviously, there is the competitive pressure that permeates everything that goes on in manufacturing and engineering design. It may, on the ground, foster irrational behavior and unsustainable solutions, but it is institutionalized in budgets, in annual and quarterly targets, in performance measurement systems, etc.

But there is also a certain ethos in the engineering approach to modeling which was nicely expressed when Pepper White’s teacher at MIT said that ‘once you know how to model things, you can model anything. It does not matter whether it’s a mechanical, fluid, thermal, chemical, electrical, or biological system. The concepts of modeling are the same’ (White, 1991, p. 122). Given such an approach, rushing in would be the norm. It would also make one inattentive to the incommensurate conceptualizations of, say mechanical, electrical, and organizational systems.

This rather rush approach to modeling is also reflected in the observable proclivity to extend the object of modeling from the factual (e.g. work processes and products) to the not so factual (e.g. contractual arrangements, trust, knowledge), as we can for example see in the conclusion Paul draws from the example of the not specified requirements: ‘... if we don’t catch the product requirement, then we will not catch the business agreement either’. And then from static object structures to evolutionary dynamics, assuming causal dependencies in people’s actions and disregarding intentions, encountering vast opportunities for disaster.

Asking for the moon

When visiting Cuba shortly after the revolution, Jean-Paul Sartre had a conversation with Fidel Castro. At one point in the conversation Castro said that the revolution would get people whatever they requested, to which Sartre raised the sensible question: ‘What if they asked for the moon?’ Castro thought for a moment and replied: ‘We may not be able to get it for them, but we would understand that they *need* it.’ (Sartre, 1961)

When workers at *Carparts*, *Newcars*, and many other enterprises are engaged in developing and tentatively pursuing a strategy of constructing computational models of the enterprise-wide design space, in order to find a way of coordinating internally and with other enterprises in global production networks, they may indeed be ‘asking for the moon’. What they do may eventually turn out to be impossible but that does not discount the obvious need.

Trying to meet the different and varying requests and requirements of their large customers in the automobile industry and at the same time trying to sort out their network of suppliers, the practitioners at *Carparts* are engaged in a very demanding exercise. The received ways of doing this, relying on a network of (passive) documents and a flow of documents is seen as increasingly inadequate. They need ‘active documents’, that is, facilities that can automate their work of keeping track of design interdependencies.

These conceptual and practical problems exemplify what Bittner (Bittner, 1965), in his brilliant essay ‘The concept of organization’, wrote about organizational rules, arguing that the sense of a organizational rule (and, *a fortiori*, a model) is relative to the practice for which it has been devised. This is reflected in his suggestion to ‘attain a grasp of the meaning of the rules as common-sense constructs from the perspective of those persons who promulgate and live with them’ [p. 251]. Interestingly he refers to the role of organizational rules in linking affiliations between entities (people, tasks, parts of a complex product, and so forth) that ‘are too remote for contingent arrangement’. Organizational rules help people link those entities into ‘coherent maps or schedules’ where ‘each link derives its meaning not so much from the specific rule that determines it, but from the entire order of which the rule itself is a part’ [p. 252]. That is, organizational rules (or models) are constructs members of a particular organizational unit or profession define in order to connect with elements that are outside the scope of their own direct influence. How these rules are understood and evoked depends on the situation, practice, and perspective of the involved actors. With this argument Bittner points to the fundamental ambiguity and openness of rules but also to their power in linking things that are remote — geographically and socially, but also conceptually.

A way to conceptualize *the specifics* of the kind of budding practice we have observed would be to discuss it under the perspective of the concept of ‘boundary objects’ (Star and Griesemer, 1989). This term was introduced and is being used to denote artifacts that, at the boundary between different local practices, facilitate loosely coupled collaboration between these communities. In the words of Bowker and Star:

‘Boundary objects are those objects that both inhabit several communities of practice and satisfy the informational requirements of each of them. Boundary objects are thus both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity

across sites. They are weakly structured in common use and become strongly structured in individual-site use.’ (Bowker and Star, 1999)

The models under tentative construction at *Carparts* were obviously conceived of as something akin to boundary objects, in as much as the models were deliberately designed to be far less detailed than the CAD models of products and parts already in use. Practitioners at *Carparts* and modelers thus made an interesting distinction between CAD drawings and product models. The product model provides a simplified view of each part of the product, hiding much of its complexity. Contextual knowledge can be added, as well as information on pending issues and on related tasks and responsibilities. That is, each model is not just a drawing; it has property sets.

Anyway, whatever their current status, what workers at *Carparts* are trying to construct goes beyond boundary objects in that the product model is obviously intended to regulate local action in a rather strong sense. This, then, poses the problem they are struggling with: they are trying to construct one integrated and overarching model for heterogeneous practices, not a family of related models representing different perspectives. In other words, what they are up against is that representations are local and temporary closures (Gerson and Star, 1986).

Now, building one integrated and overarching model may very well be the only viable approach. But it might just as well be a prejudice, if not on the part of practitioners at *Carparts*, then a prejudice on the part of developers of notations and tools of modeling. That is, perhaps a family of related models would be more feasible: more appropriate for a bottom-up process of model construction; more appropriate for involving, expressing, integrating multiple perspectives.

Modelers within MAPPER were principally aware of these problems but they were also convinced that they had the right approach to addressing them in efficient ways. A modeler at *Newcars*, for instance, demonstrated his acute awareness that a multiplicity of models is required when he talked of his approach as a war room: ‘*The idea is that for each wall of this room you have different models representing different domains. You have an expert for each of these walls and when you are in the middle, you just can give a look to all these models and try to see the connections between process and organization, product and system*’.

While reproducing the myth of an omniscient observer who can instantly see and grasp all the connections (when placed ‘in the middle’), this proposal also, albeit implicitly, demonstrates that current modeling technologies are deficient when it comes to expressing the relatedness of perspectives and thus supporting the interconnectedness of heterogeneous practices and leaves it to practitioners to figure that out themselves, as they have always done.

Existing technologies of modeling are very flexible when it comes to building models in a piecemeal fashion and then connecting them, just as they offer the flexibility of choosing different modeling approaches and notations. However, the

current modeling environments are lacking when it comes to expressing the relatedness of models from different perspectives.

That is, there is definitely a room for CSCW research to fill this gap between monolithic models and disconnected models. In fact, there is not only a room, there is a need.

Acknowledgments

The research has been supported by the European Commission (under the MAPPER project), by the Danish Research Council for Technology and Production (CosmoBiz, Grant no. 274-06-0415, the Danish Strategic Research Council (CITH, Grant No. 2106-07-0017) and by the University of Trento. We are grateful for the hospitality extended by practitioners at *Carparts* and *Newcars*. We are particularly grateful to Gianni Jacucci for his insights and comments on the modeling practices in both companies.

References

- Alexander, Christopher: *Notes on the Synthesis of Form*, Harvard University Press, Cambridge, Mass., 1964.
- Arndt, Sven W.; and Henryk Kierzkowski (eds.): *Fragmentation: New Production Patterns in the World Economy*, Oxford University Press, Oxford, 2001.
- Becker, Markus C.; and Francesco Zirpoli: 'How has the co-ordination of knowledge and competencies changed after the wave of strong engineering outsourcing in the automotive industry? The case of Italy', in *Dixième rencontre internationale du GERPISA / Tenth GERPISA international colloquium: La coordination des compétences et des connaissances dans l'industrie automobile / Co-ordinating competencies and knowledge in the auto industry, 6-8 Juin 2002, Paris, France, 2002*, pp. 1-20.
- Berger, Suzanne: *How We Compete: What Companies Around the World Are Doing to Make it in Today's Global Economy* Doubleday Publishing, New York, 2005.
- Bittner, Egon: 'The concept of organization', *Social Research*, vol. 32, no. 3, Autumn 1965, pp. 239-255.
- Bonazzi, Giuseppe; and Cristiano Antonelli: 'To make or to sell? The case of in-house outsourcing at Fiat Auto', *Organization Studies*, vol. 24, no. 4, 2003, pp. 575-594.
- Bowker, Geoffrey C.; and Susan Leigh Star: *Sorting Things Out: Classification and Its Consequences*, MIT Press, Cambridge, Mass., 1999.
- Carstensen, Anders, et al.: *Generalised AKM for Automotive Distributed Product Design, Work Package 3: Distributive Product Design in Automotive Supplier Industry*, MAPPER Consortium, 5 March 2008.
- Chandler, Alfred D., Jr.: *The Visible Hand: The Managerial Revolution in American Business*, Belknap Press, Cambridge, Mass., 1977.
- Cheng, Leonard K.; and Henryk Kierzkowski (eds.): *Global Production and Trade in East Asia*, Kluwer Academic Publishers, Boston, etc., 2001.
- Conklin, Jeff; and Michael L. Begeman: 'gIBIS: A hypertext tool for exploratory policy discussion', in I. Greif and L. A. Suchman (eds.): *CSCW'88: Proceedings of the Conference on Computer-Supported Cooperative Work, 26-28 September 1988, Portland, Oregon*, ACM Press, New York, 1988, pp. 140-152.
- Gerson, Elihu M.; and Susan Leigh Star: 'Analyzing due process in the workplace', *ACM Transactions on Office Information Systems*, vol. 4, no. 3, July 1986, pp. 257-270.

- Harris, Richard G.: 'A communication-based model of global production fragmentation', in S. W. Arndt and H. Kierzkowski (eds.): *Fragmentation: New Production Patterns in the World Economy*, Oxford University Press, Oxford, 2001, pp. 52-75.
- Jacucci, Gianni; Hilda Tellioglu; and Ina Wagner: 'Practices and cultures of knowledge management', in *CCE'07: 5th Workshop on Challenges in Collaborative Engineering, Kraków, Poland, 11-13 April 2007*, 2007, pp. 99-112.
- Levinson, Marc: *The Box: How the Shipping Container Made the World Smaller and the World Economy Bigger*, Princeton University Press, Princeton and Oxford, 2006.
- Lillehagen, Frank; and John Krogstie: 'Active knowledge models and enterprise knowledge management', in K. Kosanke, et al. (eds.): *Enterprise Inter- and Intra-organizational Integration: Building International Consensus*, Kluwer Academic Publishers, Dordrecht, 2002, pp. 91-100.
- Midler, Christophe: "'Projectification" of the firm: The Renault case', *Scandinavian Journal of Management*, vol. 11, no. 4, 1995, pp. 363-375.
- Mili, Hafedh, et al.: 'Business process modeling languages: Sorting thorough the alphabet soup' [Manuscript]. November 2003.
<<http://www.info2.uqam.ca/~mili/Recherche/publications.html>>
- Rabe, Markus; and Frank-Walter Jaekel: 'The MISSION project', in K. Kosanke, et al. (eds.): *Enterprise Inter- and Intra-organizational Integration: Building International Consensus*, Kluwer Academic Publishers, Dordrecht, 2002, pp. 235-242.
- Rittel, Horst W. J.; and Melvin M. Webber: 'Dilemmas in a general theory of planning', *Policy Sciences*, vol. 4, 1973, pp. 155-169.
- Sartre, Jean-Paul: *Sartre on Cuba*, Ballantine Books, New York, 1961.
- Schmidt, Kjeld: 'Riding a tiger, or Computer Supported Cooperative Work', in L. J. Bannon; M. Robinson; and K. Schmidt (eds.): *ECSCW'91: Proceedings of the Second European Conference on Computer-Supported Cooperative Work, 24-27 September 1991, Amsterdam*, Kluwer Academic Publishers, Dordrecht, 1991, pp. 1-16.
- Schmidt, Kjeld: 'Cooperative design: Prospects for CSCW in design', *Journal of Design Sciences and Technology*, vol. 6, no. 2, 1998, pp. 5-18.
- Smith, Adam: *An Inquiry into the Nature and Causes of the Wealth of Nations*, vol. 1-2, Printed for W. Strahan; and T. Cadell, in the Strand, London, 1776.
- Star, Susan Leigh; and James R. Griesemer: 'Institutional ecology, "translations" and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39', *Social Studies of Science*, vol. 19, 1989, pp. 387-420.
- Subrahmanian, Eswaran, et al.: 'Computer-based support for cooperative work in engineering design and manufacturing', in P. Kidd and W. Karwowski (eds.): *Advances in Agile Manufacturing. Integrating Technology, Organization and People. - Fourth International Conference on Human Aspects of Advanced Manufacturing and Hybrid Automation, Manchester, England, July 6-8, 1994*, IOS Press, Amsterdam, 1994, pp. 109-112.
- Subrahmanian, Eswaran, et al.: 'The N-Dim approach to creating design support systems', in *DETC'97: Proceedings of ASME Design Engineering Technical Conference, 14-17 September 1997, Sacramento, California*, 1997.
- Tellioglu, Hilda: 'Practicing modelling in manufacturing', in D. Dori and Y. Reich (eds.): *MBSE'09: Proceedings of the Second International Conference on Model Based Systems Engineering, 2-6 March 2009, Haifa, Israel*, 2009.
- UNCTAD: *Trade and Development Report, 2002*, United Nations Conference on Trade and Development, New York and Geneva, 2002.
- UNCTAD: *Trade and Development Report, 2005*, United Nations Conference on Trade and Development, New York and Geneva, 2005.
- White, Pepper: *The Idea Factory: Learning to Think at MIT*, Dutton, New York, 1991. - MIT Press, Cambridge, Mass., 2001 (Revised ed.).

Information Curators in an Enterprise File-Sharing Service

Michael J. Muller, David R. Millen, and Jonathan Feinberg

IBM Research, Cambridge, MA, 02142 USA

{michael_muller, david_r_millen, jdf}@us.ibm.com

Abstract. We report on a social-software file-sharing service within a large company. User-created *collections* of files were associated with increased usage of the uploaded files, especially the sharing of files from one employee to another. Employees innovated in the use of the collections features as “information curators,” an emergent lead-user role in which one employee creates named, described collections of resource for use by other employees. This role suggests new work practices and new features.

Introduction

File-sharing has been part of work practices in organizations for decades (for review, see Lee, 2003; Volda et al, 2006; Whalen et al., 2008). The advent of social software has begun to affect file-sharing activities (Reynolds et al., 2007; Schwartz, 2007), just as it has in other aspects of work in organizations and enterprises (e.g., Damianos et al., 2006; John and Seligmann, 2006; Millen et al., 2007; Muller et al., 2008; Thom-Santelli, 2008).

This paper examines an emergent behavior and role associated with file-sharing in an enterprise social software context, namely the preparation of collections of documents for use by others. We call the collectors of the documents “information curators” (see Rubel, 2008, for a similar position about bloggers as curators). Information curators are a special case of the more general role of *intermediaries* who help others to find information (Ehrlich & Cash, 1999;

Muller, 1999). Our investigation of the role of curators in enterprise file-sharing is similar to other, emergent roles in organizational social-computing contexts, such as the roles of evangelist, publisher, and community organizer described by Thom-Santelli et al. (2008) in a social-bookmarking service. Indeed, participatory Web2.0 applications tend to favor user appropriation into novel roles and work practices (Muller et al., 2005). These emergent roles and work practices can serve as lead-user descriptions (Franke et al., 2006), helping to anticipate new work practices and the designs and technologies that will be needed to support them (Kujala and Kauppinen, 2004).

This short paper is organized as follows. The next section provides a brief overview of the enterprise file-sharing service, and compares it with published reports of file-sharing and enterprise social software services. We then describe the traffic and contents of the file-sharing service, highlighting the importance of user-created collections in promoting the downloading of files from the service. The next section presents interview results from 22 of the most-active users of the collections features (i.e., users in the “curators” role). We review the work of “curators” against published requirements for file-sharing services, and we end with implications for design of social software for organizational computing.

The Cattail File-Sharing Service

The Cattail file-sharing service was originally designed as an experiment to reduce the volume of email attachments in IBM’s email service. A minimum type of functionality was thus the ability for one user to upload a file, and the ability for a second user to download a file. However, in keeping with social-software concepts, its design quickly evolved into a socially-informed venue for sharing files, networking with other users, discovering new information, and constructing aggregates of files (collections) for individual or shared purposes.

In contrast with the current research focus on peer-to-peer file-sharing networks (e.g., Christin et al., 2005; Johnson et al., 2009; Lee, 2003; see also Volda et al., 2006, for a brief survey of internet peer-to-peer systems), Cattail was designed as a single, centralized server, somewhat similar to the UD Dropbox project for a university setting (Schwartz, 2007). Because Cattail runs entirely within a protected corporate intranet, with full authentication for every user, there have been no known issues with inappropriate sharing of copyrighted materials, or with copyright-owners’ countermeasures (e.g., Cristin et al., 2005). Because Cattail does not access individual users’ own file systems (other than explicit user-initiated uploads and downloads), there have been few issues of personal data becoming visible to unintended audiences (e.g., Johnson et al., 2009).

Cattail thus shares the same type of operational intranet environment as the Apocrita peer-to-peer system (Reynolds et al., 2007), but with a centralized architecture and enhanced social-networking features. When usage data are

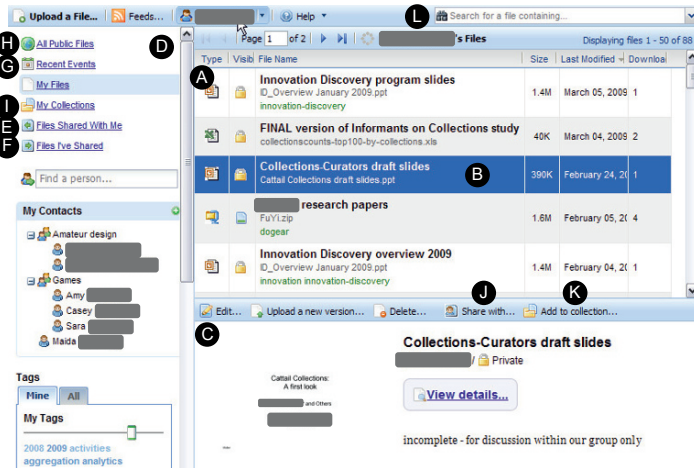


Figure 1. Cattail view of the files related to a particular user. Grey ovals obscure the names of users to protect their privacy.

available from Apocrita and for UD Dropbox, it will be interesting to compare the social interactions associated with each system.

Cattail User Experience

Figure 1 shows a view of a user's own resources in Cattail. The large window on the right (A) contains a list of the user's files. One file has been selected (B), and more information about that file is displayed in the bottom window (C).

The navigation window on the left (D) allows to display different sets of files related to the user, including "My Files" (the current view), files shared *from* other people *to* the user (E), files shared by the user *to* other people (F), recent events in the user's file-sharing network (G), a list of all public files (H), or a list of the collections of files to which the user had access (I). Clicking on the name of a collection displays the files in that collection in a manner similar to Figure 1A.

If the user selected a file (B) and requested to share it (J), the user would be prompted to supply the names of the people to whom the file was to be shared, and an option to send them a message notifying them of the share event (not illustrated). Similarly, if the user selected a file and asked to add it to a collection (K), the user would be prompted to select the collection name from a menu of collections to which s/he had access permissions. The user interface also allows to search all of the files to which the user has access permissions (L). Finally, a detailed view allows the user to add an annotation (or comment) onto the file (not illustrated).

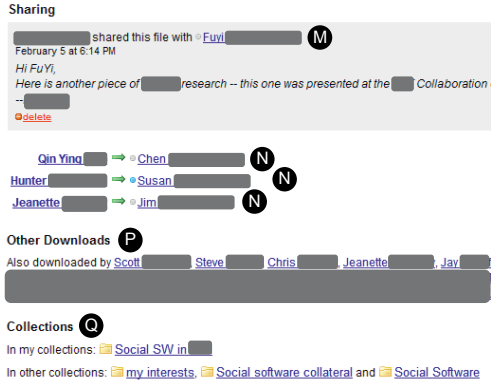


Figure 2. Detail of sharing information about one file. Cattail allows users to track what information was shared, and with whom. Cattail also shows the collections containing the file.

Cattail’s functionality reflects in part the analyses of Bellotti (1996), Volda et al. (2006), and Whalen et al. (2008). Volda et al. and Whalen et al. noted user needs to track what had been shared, and with whom – capabilities that are available in the sharing information about each file (Figure 2). Volda et al. also noted user needs to be able to notify others of new or updated material. Cattail records sharing acts by the uploader of the file, including notification messages (M); sharing acts by other users (N); a complete list of all users who have downloaded the file (P); and a list of all collections in which the file has been included (Q); these features come close to the file-history that was advocated by Whalen et al. Bellotti (1996) and Volda et al. (2006) also described a need to be able to specify a group of users, and then to share explicitly with that group; Cattail permits each collection to have a list of members (not illustrated), thus supporting this functionality. Finally, Whalen et al. advocated providing an artifact-based view and a person-based view of file-sharing. Figure 2 illustrates a file-based view, and each person’s name serves as a link to a person-based view (similar to Figure 1).

Experiences with File-Sharing in Cattail

The dataset for the quantitative analyses comprised all of the Cattail data records from its introduction on 17 May, 2007 until 10 December 2008. During that time, 15934 employees uploaded a total of 120288 files, which were accessed by a total of 85707 employees (including the 15934 uploaders), who collectively performed 728509 download events. Employees spanned 81 countries, and were employed in a diversity of organizations including product development, sales, marketing, planning, internal operations, and research. Table 1 provides a high-level summary of system usage during the study period.

Resources		Principal Actions		Users in roles	
Files	120288	Upload	120288	Total	88270
Collections	12461	Download	728509	Uploaders	15934
Annotations	8828	Share	107341	Downloaders	85707
		Collect	79823	Sharers	12584
		Annotate	9494	Collectors	5444
				Annotators	3884

Table 1. Cattail resources, actions, and user roles from 17 May 2007 through 10 December 2008.

We were initially interested to predict which files would be downloaded – i.e., Cattail exists to serve files to users through the download operation. While distinct download operations occurred 728509 times, only 89956 unique files (just under 75% of all files) were downloaded. Therefore, despite the purpose of Cattail to share files, more than 25% of the stored files were never downloaded.

Through a multiple regression analysis, we found a higher number of downloads for files that had been shared ($F_{(1,127287)}=7696$, $p<.001$), collected ($F_{(1,127287)}=5501$, $p<.001$), and/or annotated ($F_{(1,127287)}=5104$, $p<.001$).

It seemed obvious to us that sharing a file should increase the number of downloads of that file. We therefore did not focus on the sharing operation.

The work practice of collecting files was more interesting to us. As shown in Table 1, 5444 users created 12461 collections of files through 79823 distinct operations of adding a file to a collection. A total of 60476 unique files (50%) appeared in at least one collection (ten times the number of files with annotations). The phenomenon of collecting affected half of all files. We also found that collections were only marginally associated with refinding of one's own files, ($F_{(1,127287)}=4.6$, $p<.05$), while they were highly associated with downloading files that had been uploaded by others ($F_{(1,127287)}=4910$, $p<.001$).

Qualitative Exploration of Collections

To gain insight into the phenomena associated with collections, we undertook an interview study. We identified the 100 most-frequent users of the make-collection feature, including employees from 22 of the 81 countries in the initial dataset of 15934 users. We selected 22 people to interview, using heuristics of trying to maximize diversity in number of countries (16), and to balance women (36%) and men (64%). Informants held a diverse array of job titles, including sales, consulting, business operations, solutions architecture, product management, internal communications, and design.

To reduce costs, and to accommodate large differences in timezones, interviews were conducted via one-hour instant messaging sessions. Interviews were semi-structured, guided by the following topics: motivation for creating collections; users of the informant's collections; presentation/interpretation of the

informant's collections to their users; use of other people's collections. Interview results were coded by a single analyst, through five coding iterations. For this initial analysis, we focused largely on the answers to the above questions.

Informants used collections for individual and shared work: *"easier for myself and others to find the content again... a more active way of sharing"* (Informant 11, internal communications, Austria); *"a knowledge package.... to summarize and to create mixed knowledges [with others]"* (18, consultant, Turkey); and *"group them in a neat bundle to shove it at people"* (19, project management, UK) as *"a single focal point of entry"* (I13, enablement, Spain).

Many informants had a semi-structured approach to using collections: *"1. organizing files... 2. finding files I use most often (either my own or those of the people i [work with]... 3. sharing files..."* (I22, sales, USA). In some cases, the collection itself was highly structured: *"collection will contain a master report, and an updated report for every member of my department.... 31 people are able to download reports. 20 people are able to update that report.... save it on cattaill so I open a new collection every week..."* (I2, supply chain specialist, Mexico).

Several informants reported using multiple collections for multiple audiences: *"organize by clients or projects... by [human capital management] topic (i.e., workforce planning, hr Business Intelligence) or by project (i.e: customer XX project definition)..."* I17, consultant, Italy); *"Rather than share each [article], I just shared the collection... [project1] internal initiative... [project2] Community calls... [project3] share free ebooks from industry..."* (I20, designer, Canada); *"my team... my boss and his team... virtual development teams... brand executives and their orgs... [world-wide] enablement folks..."* (I22, sales, USA).

Information Curators

Informants created collections with specific intentions for their use, and detailed concern for their audiences: *"regular collections with manually selected / curated resources.... trying to help people (and myself!) make sense of the files that are available.... putting together a collection and deciding what goes into it... and if they are different from the ones I've seen before then I add them to my collection..."* (I15, enterprise 2.0 evangelist, Canada); *"a kind of editor, you share you own and other useful info via collections"* (I18, sales, Finland); *"put some structure around the content I collect/create around my topic... what is good for me is good for my readers ☺"* (I19, product manager, France).

People in this emergent role were concerned to describe or frame their collections for discovery and use by others: *"very short descriptions... the intent of the collection – so I can keep the collection name really short!"* (I9, project management, UK); *"sometimes I used the [descriptive field] to link to other related content [cross] reference"* (I19, product manager, France); and *"i asked everyone to use the naming convention, and I enforced it"* (I22, sales, USA).

These curated or edited collections were intended for both current and future use: “*audit [can] go in on a monthly basis so that they can test to see if the necessary billing approvals exist.*” (14, business operations, UK); “*an asset for the future opportunities about the client*” (18, consultant, Turkey) and “*It’s a fail safe if I was knocked down by a bus!*” (14, business operations, UK).

Conclusions: Implications for Design

We have shown through quantitative analyses that collections are strongly associated with the use of uploaded files. Collections are particularly important in promoting use by other users – i.e., downloading by a user who had not herself uploaded that file. Collections appear to be key aspects of making a file-sharing service work in a social software environment.

Qualitative analyses illustrated how employees have adopted and adapted collections into their work practices with existing teams and work domains. Qualitative analyses also showed an emergent new role, that of the “information curator,” who prepares assemblies of materials for known audiences, and who presents or interprets those materials to those audiences. Curators’ collections have lasting value to their audiences and, potentially, to their organizations.

Curators may provide a view into the future of file-sharing, similarly to other “lead user” roles that have helped to define new practices, new opportunities, and new features (Franke et al., 2006; Kujala and Kauppinen, 2004; Thom-Santelli et al., 2008b). Unlike the blogging “digital curators” proposed by Rubel (2008) and others, the information curators in our study often knew their audiences, and collected files to match specific audience needs. Curators have used the social attributes of Cattail to address some of the needs outlined in previous research, such as the ability to share easily with a known audience, to create and use views based on both artifacts (and now *collections* of artifacts) and on persons (Bellotti, 1996; Volda et al., 2006; Whalen et al., 2008).

However, informants also noted gaps in the functionality. While a person-centered view is useful, there are no *group*-centered views, e.g. of the downloads or other actions of the *audience* of a collection. Thus, audience analysis and audience development remain major challenges (see also Thom-Santelli et al., 2008). While a historical view of the actions related to a *single file* is valuable, there is no means for summarizing the history of a *collection* of files. Curators must work hard to understand if their collections are being used, and especially if each collection is providing value *as a collection*. While tagging and annotating are available to clarify the meaning and significance of individual files, there are no similar capabilities to present, discuss, and co-create the meaning and significance of a collection as an aggregate. We will explore these types of new features, and we will be eager to see whether curators emerge as lead users in other organizational file-sharing projects (Reynolds et al., 2007; Schwartz, 2007).

References

- Bellotti, V. (1996). 'What you don't know can hurt you: Privacy in collaborative computing.' *Proc BCS HCI '96*. Springer-Verlag, London, UK, January 1996 241-261.
- Christin, N., Weigent, A., & Chuang, J. (2005). 'Content availability, pollution, and poisoning in file sharing peer-to-peer networks.' *Proc EC'05*, ACM Press, Vancouver, BC, Canada, June 2005, 68-77.
- Damianos, L., Griffith, J., & Cuomo, D. (2006). 'Onomi: Social Bookmarking on a Corporate Intranet.' Position paper in WWW 2006 Tagging Workshop, Edinburgh, Scotland, UK, May 2006, <http://www.semanticmetadata.net/hosted/taggingws-www2006-files/28.pdf>.
- Ehrlich, K., & Cash, D. (1999). 'The invisible world of intermediaries: A cautionary tale,' *Computer Supported Cooperative Work*, vol. 8 no.1-2, p.147-167.
- Franke, N., von Hippel, E., & Schreier, M. (2006). 'Finding commercially attractive user innovations: A test of lead-user theory.' *Jour. Prod. Innov. Mgmt.*, vol. 23, no. 4, 301-315.
- John, A., & Seligmann, D. (2006). 'Collaborative tagging and expertise in the enterprise.' Position paper in WWW 2006 Tagging Workshop, Edinburgh, Scotland, UK, May 2006, <http://www.semanticmetadata.net/hosted/taggingws-www2006-files/26.pdf>
- Johnson, M.E., McGuire, D., & Willey, N.D. (2009). 'Why file sharing networks are dangerous?' *Communications of the ACM*, vol. 52, no. 2, February 2009, 134-138.
- Kujala, S., & Kauppinen, M. (2004). 'Identifying and selecting users for user-centered design.' *Proc NordCHI '04*, ACM Press, Tampere, Finland, October, 2004, 297-303.
- Lee, J. (2003). 'An end-user perspective on file-sharing systems.' *Communications of the ACM*, vol. 46, no. 2, February 2003, 49-53.
- Millen, D.R., Yang, M., Whittaker, S., & Feinberg, J. (2007). 'Social bookmarking and exploratory search.' *Proc ECSCW 2007*, Springer-Verlag, Limerick, Ireland, Sep. 2007, 21-40.
- Muller, M.J. (1999). 'Invisible work of telephone operators: An ethnocritical analysis.' *Computer-Supported Cooperative Work*, vol. 8, no. 1-2, 31-61.
- Muller, M.J., Geyer, W., Brownholtz, B., Dugan, C., Millen, D.R., and Wilcox, E. (2007). 'Tag-based metonymic search in an activity-centric aggregation service.' *Proc ECSCW 2007*, Springer-Verlag, Limerick, Ireland, September 2007, 179-198.
- Muller, M.J., Minassian, S.O., Geyer, W., Millen, D.R., Brownholtz, E., & Wilcox, E. (2005). 'Studying appropriation in activity-centric collaboration.' Position paper at ECSCW 2005 workshop, *Supporting appropriation work*.
- Reynolds, J.J., McLeod, R., & Mahmoud, Q.H. (2007). 'Apocrita: A distributed peer-to-peer file sharing system for intranets.' *Proc ACMSE'07*, ACM Press, Winston-Salem, NC, USA, March 2007, 174-178.
- Rubel, S. (2008). 'The digital curator in your future.' *Micropersuasion* <http://www.micropersuasion.com/2008/02/the-digital-cur.html>, 6 Feb. 2008.
- Schwartz, A. (2007). 'UD Dropbox 2.0: Collaboration magic.' *Proc SIGUCCS'07*, ACM Press, Orlando, FL, USA, October 2007, 305-309.
- Thom-Santelli, J., Muller, M., & Millen, D.R. (2008). 'Social tagging roles: Publishers, evangelists, leaders.' *Proc CHI 2008*, ACM Press, Florence, IT, April 2008, 1041-1044.
- Voida, S., Edward, W.K., Newman, M.W., Grinter, R.E., & Ducheneault, N. (2006). 'Share and share alike: Exploring the user interface affordances of file sharing.' *Proc CHI 2006*, ACM Press, Montréal, QU, Canada, April 2006, 221-230.
- Whalen, T., Toms, E.G., & Blustein, J. (2008). 'Information displays for managing shared files.' *Proc CHIMIT'08*, ACM Press, San Diego, CA, USA, November 2008.

Index of Authors

Aarhus, Rikke.....	63	Martin, David.....	171
Anderson, Theresa Dirndorfer.....	21	Millen, David.....	143, 403
Ashkanasy, Shawn.....	105	Muller, Michael.....	143, 403
Ballegaard, Stinne Aaløkke.....	63	Nett, Bernhard.....	123
Bietz, Matthew J.....	243	Nijholt, Anton.....	151
Boden, Alexander.....	123	O'Neill, Jacki.....	171
Borchorst, Nikolaj Gandrup.....	1	Patterson, Donald.....	303
Boujut, JF.....	191	Piper, Anne Marie.....	283
Brosch, Petra.....	207	Randall, Dave.....	171
Bødker, Susanne.....	1	Rosa, Maryanne.....	215
Cabitza, Federico.....	85	Schmidt, Kjeld.....	223, 383
Carpendale, Sheelagh.....	43	Schümmer, Till.....	263
Costa, Jean.....	215	Seidl, Martina.....	207
Darses, Françoise.....	191	Shanks, Graeme.....	105
Davis, Hilary.....	105	Simone, Carla.....	85
de Souza, Cleidson.....	215	Tang, Anthony.....	343
Ding, Xianghua.....	303	Tang, Charlotte.....	43
Ducheneaut, Nicolas.....	323	Tellioglu, Hilda.....	263, 383
Dugan, Casey.....	143	Thom-Santelli, Jennifer.....	143
Feinberg, Jonathan.....	403	Trecanno, Pedro.....	215
Freyne, Jill.....	143	van der Veer, Gerrit.....	151
Goto, Crys.....	215	Vetere, Frank.....	105
Guibert, Sylvie.....	191	Vyas, Dhaval.....	151
Gutwin, Carl.....	343	Wadley, Greg.....	323
Haake, Jörg M.....	263	Wagner, Ina.....	383
Hansen, Thomas Riisgaard.....	63	Wieland, Konrad.....	207
Heyer, Clint.....	363	Wimmer, Manuel.....	207
Heylen, Dirk.....	151	Wong, Nelson.....	343
Hollan, James D.....	283	Wulf, Volker.....	123
Langer, Philip.....	207	Zander, Pär-Ola.....	1
Lee, Charlotte P.....	243	Zorzato, Giovanni.....	85
Light, Ann.....	21		
Livingston, Ian.....	343		
Mandryk, Regan.....	343		