Chapter 12
Two Thousand Points of Interaction: Augmenting Paper Notes for a Distributed User Experience

Gunnar Harboe, Gelek Doksam, Lukas Keller, and Elaine M. Huang

Abstract We present two early prototypes of a system that couples an augmented wall of paper notes with multiple handheld devices in order to support the process of affinity diagramming. Our system allows multiple users to work together simultaneously, freely interacting with potentially thousands of physical notes directly, and with a coupled digital representation of the same notes via a smart phone, tablet or PC. We propose the affinity diagramming process as a use-case well suited for distributed user interfaces.

12.1 Introduction

Some activities are not well suited to be performed within the constraints of a traditional computer interface and context of use, and are therefore still accomplished mainly by "manual" (i.e., non-computerized) methods. Distributed user interfaces (DUIs) potentially offer new ways to accomplish these activities. In this paper we examine one such activity, affinity diagramming, and present two prototypes that create a distributed user experience using augmentation of physical objects (up to thousands of paper notes) as well as digital devices.

In the following section we explain the affinity diagramming method as it is carried out today, and discuss why paper-based processes are preferred over desktop computer alternatives. Identifying contributing factors such as support for group interaction and collaboration, tangibility, spatial awareness and interaction richness, we argue that this suggests a potential role for DUI solutions.

G. Harboe (✉) • G. Doksam • L. Keller • E.M. Huang
University of Zurich, Zurich, Switzerland
e-mail: harboe@ifi.uzh.ch; gelek.doksam@uzh.ch; lukas.keller@uzh.ch; huang@ifi.uzh.ch

M.D. Lozano et al. (eds.), *Distributed User Interfaces: Usability and Collaboration*, Human–Computer Interaction Series, DOI 10.1007/978-1-4471-5499-0_12,

We then give an overview of related work, and argue that an approach based on large touch displays presents a barrier to use. Instead we propose a hybrid approach of augmenting the physical diagram and coupling it with additional interactions on standard mobile devices. We describe our design concept for a comprehensive solution, as well as two currently implemented prototypes towards that vision. We conclude by considering the more general applications for this type of distributed user interface.

12.2 Background on Affinity Diagramming

Loose paper notes arranged on a surface are commonly used to support a variety of tasks that involve creativity, design, planning, and organizing or analyzing information. Examples include film storyboards, scrum boards in agile software development, brainstorming, studying for an exam, or planning books on note cards.

One particular process of arranging paper notes on a surface is called *affinity diagramming* (or sometimes the KJ method) [1]. In affinity diagramming, items (e.g. interview quotes, observations) are written down on separate notes, which are placed one by one on a surface (a table, or more often a wall). As the notes are placed, and in moving them around later, they are clustered based on their *affinity*: their similarity or relevance to a shared topic. This leads to the creation of groups, which are then labeled and clustered in a similar way. The process repeats until the highest level has only a few groups. In this way, affinity diagramming is a bottom-up approach to organizing unstructured data into hierarchical categories.

Affinity diagramming is used to several different ends in a number of disciplines, including CSCW and HCI as well as design, business and anthropology, and different variations of the method exist. We will focus on affinity diagramming to support the analysis of qualitative data, or the *affinity analysis* technique, as it is used within user-centered research and design.

12.2.1 Affinity Analysis

Many frequently used data-gathering methods for user-centered design and research (e.g., interviews, probes, participant observations, and focus groups) generate large amounts of unstructured qualitative data. In order to make sense of the information, the teams involved need to immerse themselves in the data, interpret it, and find ways to structure it. Affinity diagramming is one way to do this; other common approaches include coding the data, quantifying it in various ways, and examining and interpreting it in an informal fashion. One of the often-cited advantages of affinity analysis is that it is well suited to collaboration, enabling a team to efficiently generate a shared analysis together (Fig. 12.1) [1–3].

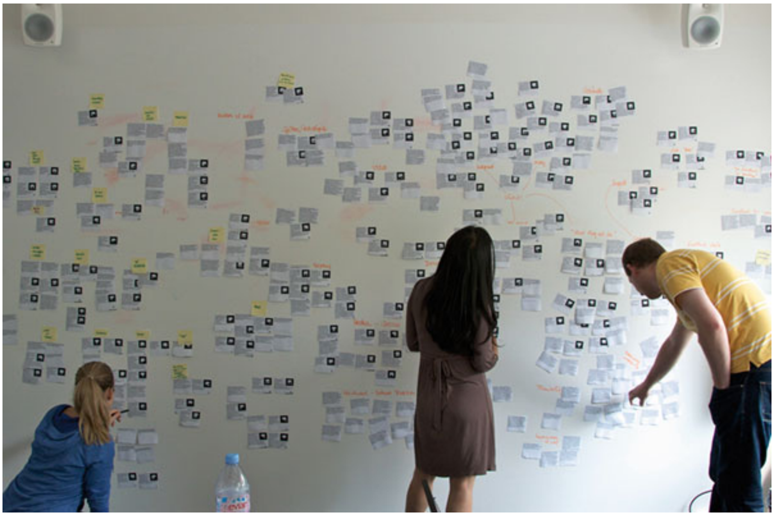

Fig. 12.1 Collaborating on an affinity diagram

A small affinity diagram for data analysis may consist of no more than a hundred notes, and a team of two may complete it in a few hours. At the other end of the scale, large affinity diagrams may contain up to 2000 notes, and teams of six to eight people sometimes work on them for weeks [4].

While a number of special- and general-purpose desktop applications could in principle be used to perform affinity analyses, previous research has found that users of the method still predominantly prefer physical paper notes [4–6] . Among the reasons given for this is that the computer tools do not have the immediacy, portability or flexibility of paper, that they do not adequately support collaboration through simultaneous interaction, that they do not offer the immersion, spatial awareness and tangibility of paper notes on a wall, and they limit the interaction richness to a single mouse cursor, hindering manipulation of multiple notes.

At the same time, the analog process has some definite drawbacks. The whole analysis is embodied in a single physical artifact (the diagram itself), which must typically remain in place throughout the process, and which is easily degraded (e.g., by notes falling down from their assigned places on the wall). As the amount of data increases, the affinity diagram becomes increasingly unwieldy, taking up a great deal of space, making it hard to find particular notes or get an overview of the complete structure, and becoming time-consuming to rearrange.

Once the process is complete, the results of the analysis usually need to be computerized. Also, since the diagram can only represent one way to organize the data, bringing out other dimensions requires the disassembly of the diagram; there is no straightforward way to store multiple versions or maintain a record or history of the process.

12.3 Related Work

In response to the challenges of paper-based affinity diagramming and similar tasks, and mindful of the limitations of desktop applications, several previous authors have built systems based on large interactive displays.

Judge et al. [6] created an all-digital affinity system using a wall-sized multiple-display environment. *The Designers' Environment* [7] is similarly an all-digital system using tablets and an interactive tabletop.

The Designer's Outpost, in contrast, combines an interactive whiteboard with physical paper notes [3] to support design brainstorming, and Geyer et al. [5] use a vertical display and interactive tabletop to capture and create digital copies of physical notes.

While these systems support collaboration, direct manipulation and spatial awareness better than desktop tools, they have their own limitations. The scale of affinity diagram they can comfortably support remains limited, they require sizable investments in dedicated hardware, and to take advantage of any of their features requires an upfront commitment to use them throughout the process, breaking with established practices.

We take a different approach, and while we don't claim our solution is better, we offer another option on the continuum between a wholly analog and a wholly computerized process.

12.4 Our Basis for Design

Instead of imitating the affordances of paper on computer screens, our idea was to support the existing paper-based processes by digitally augmenting the physical notes. In addition to the literature, we were inspired by personal experiences with affinity diagramming, which the first and last author had used in previous projects.

Not only had we faced a number of the problems described above, we had also seen attempts to overcome them. One of these solutions involved having all the notes in a file on the computer, printing them out with an identifying bar code on each, and using a barcode reader to rapidly capture the structure of the diagram upon completion. Another consisted of manually photographing the entire diagram in sections after each session.

Interviews were also conducted with researchers and practitioners who use affinity diagramming and related techniques in their work (nine so far), to gather a wider range of experiences and perspectives, and for the benefit of team members who were less familiar with the affinity method. The first few of these interviews are reported in a previous publication [4].

Based on this information, we came up with the general design concept and decided which use-cases to focus on in the initial prototypes.

12.5 Concept

It was an important design goal to keep the barriers to use as low as possible, so as not to lose the simplicity that is one of the main advantages of paper. Therefore, the design concept relies on off-the-shelf hardware that is likely to already be present in many offices and labs. It is designed to be modular, so that users may include only components they have available and wish to use, and flexible, so that users can move seamlessly between the augmented system and the traditional, plain-paper process. The system will make best-effort attempts to provide as much functionality as possible given the available resources and data.

Figure 12.2 shows an overview of the system in sketch form. The essential element of the system is capturing the content and position of each note on the wall. In one feasible implementation, this is achieved by tagging each note with a unique ID, stored in a 2D barcode (we use QR codes). A camera pointed at the wall can then recognize the barcodes. If the notes are prepared in advance on a computer, the content of each note may simply be looked up in a database using the ID. If the database does not contain the content, the best available picture of the note taken by the camera can be used. OCR or handwriting recognition may potentially be run on this picture in order to convert the content into text form.

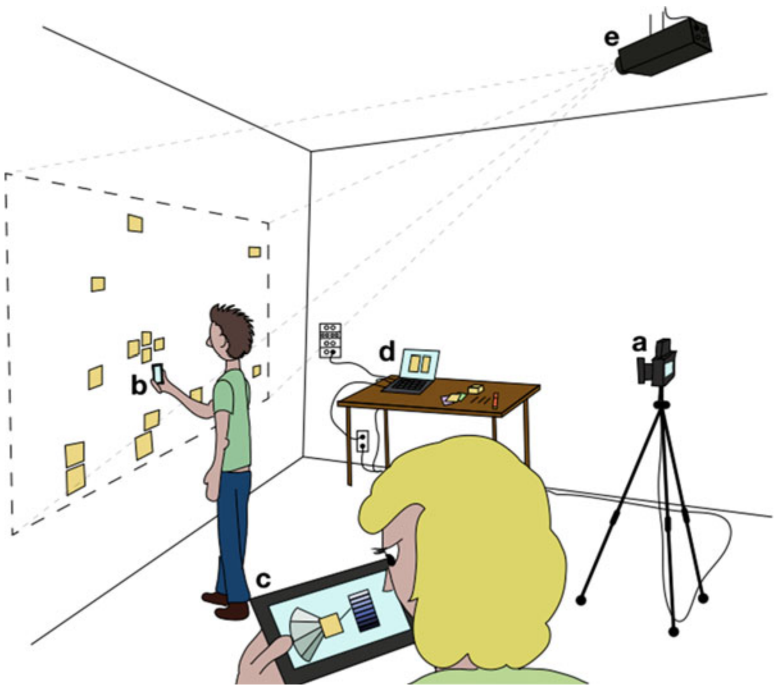

Fig. 12.2 Sketch of the augmented paper affinity diagram concept

The system takes advantage of two types of cameras; a stationary, high-resolution camera mounted on a tripod (a) to capture the whole wall, and the handheld cameras in mobile devices such as smart phones (b) and tablets (c). The stationary camera is connected to and controlled by a computer (d), and snaps photos at regular intervals during the affinity sessions, thus providing an incomplete (due to occlusions and barcode recognition failures) record of the locations of each note throughout the process.

If a stationary camera is not used, users of the system may instead photograph the affinity diagram manually, piece by piece. The system should stitch the photos together into one coherent representation, though in this case the temporal resolution of the digital copy will be lower. Data captured from the viewfinder when using the mobile clients to interact with the system can also be used to fill in and update the digital copy.

By moving the physical notes around, users indirectly manipulate the digital model of the diagram; in this way, the paper notes become points of interaction with the system. If they want to interact directly with other aspects of the system (for example to perform a search for a note) or view information that is not visible in the physical diagram, they can use clients running on smart phones, tablets and PCs, as described below.

Finally, a projector (e), connected to a computer running a system client, may be used to overlay output directly on the wall and onto the physical affinity diagram.

12.5.1 Use-Cases

The augmented affinity diagram can support many different tasks. To guide our early design and prototyping efforts, we decided to focus on some simple use-cases: (1) Locating a note on the wall based on its content; (2) Viewing the distribution of a set of notes defined by a certain characteristic; (3) Consulting additional information about a note.

12.6 Implementation

We have created prototypes of the key components of the system, allowing a user to interact through an Android smart phone or tablet client, using the projector to display output and getting input from a stationary camera.

We use a simple client–server architecture, where Android apps on the smart phone and tablet (and what is currently a Java client on the computer connected to the projector and camera) communicate over HTTP with a backend server, which handles any heavy processing and maintains the digital model of the affinity diagram in a database.

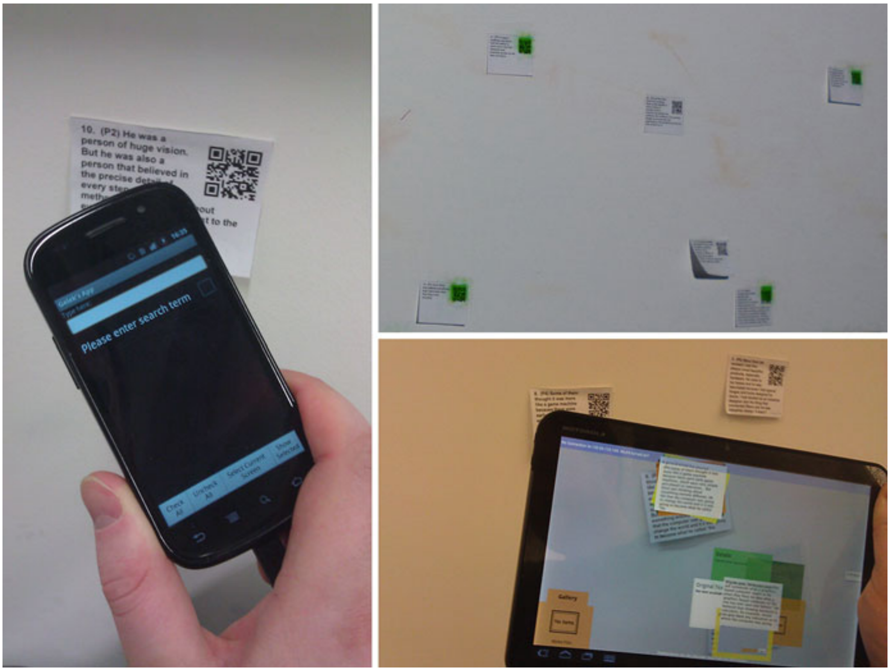

Fig. 12.3 Smart phone and tablet clients; *top right*, notes matching a search being *highlighted in green* by the projector

As part of the initial setup, the projector and stationary camera need to be calibrated to each other, in order to know where to display information that should be overlaid onto the wall. We assume the wall to be a flat surface, and so we simply display 2D barcode patterns at the corners of the projected area. By recognizing these barcodes in a photo from the camera, we can calculate the transformation between the camera view and the projector view and position overlays accordingly. Of course, if the stationary camera is moved, the projector must be recalibrated.

The current smart phone client (Fig. 12.3) offers the ability to search for notes based on their content (use-case 1). A simple search field is used to enter queries, returning matching notes as you type. Notes selected from the results are highlighted on the wall. Through the same interface, a user may also search for categories of notes, such as all notes from a certain interview, and highlight all of them at once (use-case 2).

The current tablet client, meanwhile, is designed for viewing additional information about a note, such as metadata (e.g. the participant profile, or information about the interview or observation session the note came from), photos associated with the note, and direct access to the data in raw form (e.g. the audio or video clip that the note is based on, an unedited transcript, etc.).

The interaction is again based on a magic lens model, where the user holds the tablet up to a note she is interested in more information about, while the screen displays the viewfinder view. When a barcode is recognized, the app overlays an interactive digital version of the note on top of the image of the physical note. The view can then be frozen so that the user may interact with the digital note without having to keep pointing the tablet to the physical note.

The additional information about the note is displayed as a stack of notes hidden under the note itself, which can be expanded out to sit around it and moved around freely in the neighborhood using touch gestures. Color coding is used to distinguish different types of information.

12.6.1 Limitations

These early prototypes have been created both as technology tests and to test the designs. Both aspects still suffer from many rough edges. On the technology side, the biggest challenge is the barcode recognition. While the resolution of the camera is fully able to resolve the barcode pattern, variations in brightness across the image mean that many QR codes fail to be recognized under anything other than optimal lighting conditions.

Small-scale usability tests of the prototypes have shown that participants are able to complete the intended tasks with only minor difficulties, and pointed towards some possible improvements for future iterations. The tests also indicated that the coupling of the physical notes with the digital augmentations is compelling and easily grasped.

Neither the designs nor the implementation are yet sufficiently mature to test in a realistic context, so we cannot say how well they will fit into the affinity diagramming workflow and group interaction.

12.7 Discussion

The augmented affinity diagram offers multi-user interaction distributed over a number of digital devices. However, we believe the more significant aspect of the work for distributed user interfaces is the augmentation and instrumentation of mundane objects, here in the form of paper notes, turning each of them into a point of interaction with the system. This form of ubiquitous computing is a radical break from traditional computer UIs, and here we are barely scratching the surface of its potential and what it means.

Affinity diagrams lend themselves well to this kind of augmentation, as the notes can be tagged to be easy to recognize, and the activity is relatively straightforward to model. However, as different types of sensors become commonplace and computer sensing improves, the same paradigm extends to many other activities that involve distributed interaction with physical objects, such as browsing a book shelf or a shopping aisle, cooking, tidying up and cleaning, and many types of office work. These and others will provide many interesting problems for DUIs to address.

Acknowledgements Thanks to Ioana Ilea and Jonas Minke for their contributions in the early stages of the project.

References

1. Beyer, H., & Holtzblatt, K. (1998). *Contextual design: Defining customer-centered systems*. San Francisco: Morgan Kaufmann.
2. Judge, T. K., Pyla, P. S., McCrickard, D. S., et al. (2008). *Studying group decision making in affinity diagramming*. Technical report published by Virginia Tech: TR-08-16. http://vtechworks.lib.vt.edu/handle/10919/19418.
3. Klemmer, S. R., Newman, M. W., Farrell, R., et al. (2001). The designers' outpost: A tangible interface for collaborative web site. *Proceedings of the 14th Annual ACM Symposium on User Interface Software and Technology* (pp. 1–10).
4. Harboe, G., Minke, J., Ilea, I., & Huang, E. M. (2012). Computer support for collaborative data analysis: Augmenting paper affinity diagrams. *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work.*
5. Geyer, F., Pfeil, U., Höchtl, A., et al. (2011). Designing reality-based interfaces for creative group work. *Proceedings of the 8th ACM Conference on Creativity and Cognition* (pp. 165–174). ACM, New York.
6. Judge, T. K., Pyla, P. S., McCrickard, D. S., & Harrison, S. (2008). Using multiple display environments for affinity diagramming. Beyond the laboratory: Supporting authentic collaboration with multiple displays. In *ACM Conference on Computer Supported Cooperative Work (CSCW 2008)*, San Diego.
7. Tse, E., Greenberg, S., Shen, C., et al. (2008). Exploring true multi-user multimodal interaction over a digital table. *Proceedings of the 7th ACM Conference on Designing Interactive Systems* (pp. 109–118). ACM, New York.