

# Chapter 4

## Patterns of Work: A Pragmatic Approach

K. Kashimura, Y. Hara, N. Ikeya, and David Randall

### 4.1 Introduction

As was suggested in Chap. 1, we need to ‘... accumulate and synthesize knowledge about such social domains from case studies to be able to anticipate the use and behavioral impact of new designs’. In turn, Beringer suggests, we need to go about, ‘extracting key findings and summarizing key insights [so that] they can become a reusable set of foundational insights about target domains’. How we might do this, however, is a somewhat intractable problem. The past 20 years and more has, without question, seen a significant shift in the way in which data relating to design problems is collected and analysed. One of the most significant aspects of this has been the ‘turn to the social’ often associated with the deployment of ethnographic practices for design-related purposes.

Despite the undoubted gains that come from this move away from the ‘formalities’ of business processes towards a ‘real world, real time’ analysis, we have, as yet, little insight into the way ethnographic practices and design work get mobilised in commercial practice in a timely and generalisable way. A ‘turn’ which incorporates a view of the subtle, interactive processes by which work actually gets done and the contextual character of breakdown and repair work has not as yet been shown to translate conveniently into general design features. That is, there are few studies which show how ethnographic studies might be used and deployed by design teams

---

K. Kashimura • Y. Hara  
Hitachi Design Division, Tokyo, Japan

N. Ikeya  
Keio University, Tokyo, Japan

D. Randall (✉)  
School of Media and Information, University of Siegen, Siegen, Germany  
e-mail: [D.Randall@mmu.ac.uk](mailto:D.Randall@mmu.ac.uk)

(although we should acknowledge that there are systematic attempts to specify the relation, for instance, Beyer and Holzblatt 1998), how their value is constructed and what problems with the record need to be dealt with. In addition, as we will argue below, since it is by no means clear that the highly detailed and radically contextual view implicated in (some) ethnographic work might be made to fit with some level of ‘generic’ design, some examination of how a design team might deal with this would be useful. This chapter examines the on-going work of a team of researchers and designers in the UX group at the Hitachi Design Division based in Tokyo, Japan. The Hitachi team has used ethnographic techniques for a significant period of time now, largely in the study of ‘infrastructural’ projects of one kind or another mining, power plants, train repair and maintenance and so on. Whilst considering ethnographic investigation to be valuable, the team believes that results can be made even more useful through a process of comparison. In this chapter, we consider how, in the first year of a 2-year project, this has been done. To be clear, there is a substantial existing literature on the relationship between ethnography and design (see Crabtree et al. 2012 for a recent example) which has covered a range of themes including the issues of participation and co-creation. In addition, there is an existing literature on theoretical devices which might help in the analysis of ethnographic work (see, e.g., Hutchins 1995; Kaptelinen and Nardi 2006 in the context of CSCW and other technology research; Glaser and Strauss 1967; Strauss and Corbin 1997; Charmaz 2006 in a more general context). The problem we describe, however, is somewhat different. Where Strauss and Corbin, for instance, describe the constant comparative method as a means to provide for comparisons across settings but within a domain and where the theoretical positions we mention provide frameworks for analysing any and all settings, albeit at a high level, our concern lies in comparisons across different settings and domains, but in which we may nevertheless find points of comparison and moreover comparisons which might have a design relevance.

## 4.2 The Setting

The design team’s work has, for some years, been substantially ethnographic in its orientation, and studies have been conducted in a large variety of domains. Nevertheless, as the company has refocused its efforts towards ‘quality of service’ issues in complex organisational and inter-organisational settings, we have identified various kinds of maintenance work as a core topic. The group has largely oriented over a period of time to some version of ‘contextual design’ (Beyer and Holzblatt, *ibid*). Despite this, the group has expressed certain concerns about this mode of operation, in the main a function of the sheer number of enquiries being undertaken and, at the same time, the range of design possibilities under consideration. Put simply, at any one time there might be between 6 and 12 people doing ethnographic work in different areas (normally in ones and twos) and reporting to a moveable feast in respect of design teams. Fieldworkers can be, and are, deployed at relatively short

notice into domains of which they may have no prior knowledge. They typically have had little background in ethnographic research, although a number have a background in 'usability' testing of one kind or another. The concerns of the group have effectively centred on three related difficulties. Firstly, and it will come as no surprise to experienced ethnographers, there are doubts about the capacity of designers to deal with significant amounts of ethnographic data either in relation to written records or in meetings. Although there will be more than one reason for this, in our case this mainly had to do with the structuring of data so as to reflect what designers might construe as their most significant problems. Secondly, and given the large amount of work being done in different domains, there was a concern for consistency and focus in respect of reporting. That is, how reporting might be done in such a way that ethnographic reports have a recognisable character from one setting to another was seen to be a problem, in that designers wanted ethnographic data to take a typical and consistent form and one, moreover, which might lead to relatively clear problem statements. Thirdly, there was a concern for time taken in the conduct of ethnographic enquiry. Designers, as is well known, cannot wait whilst long-term studies are completed. In the 'real world', there are also difficulties with the concurrent model described by Hughes et al. (1992, 1993). We should perhaps stress here what they are, because it has important consequences for the way in which our work progressed. The way in which ethnographic results might feed into design has important commercial implications. The team needs to demonstrate that it can provide useful and significant results both within and across a number of settings and, moreover, represent those results to a heterogeneous group of designers. Design scenarios might include, for instance, information system design, work and procedural design, augmented reality, mobile devices, video conferencing and data sharing, distributed database provision and so on. The large number of people involved at one time or another means that long-term and 'personal' collaborations of the kind often found in academic environments are difficult to pursue. Although there is a preference for face-to-face meetings between ethnographers and designers, this is not always possible. When it is, marshalling the vast range of data that ethnographers have collected and translating it into a language that is relevant and useful for design have proven less than straightforward. The ethnographic record itself, in the unstructured form that social scientists are used to, has proven unpopular. Various 'storyboarding' efforts have been undertaken, with limited success. We should point out here that, although this is not the topic of this chapter, a series of interviews with system designers of one kind or another (who described themselves variously as 'product', 'system' or 'future' designers) were conducted whilst the work was progressing, with a view to establishing what their informational and representational needs might be. One of the outcomes of these interviews was the recognition that a simplified and relatively standardised way of representing results was required if they were to derive the kinds of problem statement they felt they needed.

As a result, and in recognition of the need to develop some kind of reporting structure, a small group along with academics from Japan and Europe were convened to examine ways of dealing with the somewhat vaguely glimpsed problems

mentioned above. At the outset, and in the first instance done largely through email and Skype communications, an attempt was made to discuss initial moves before contracts were signed and analytic work began. A number of Skype meetings took place in which discussions predicated on the possibility of using ‘patterns’ as an analytic device for producing a standard presentation format were the main focus, followed by the first of two intensive workshops, each lasting 5 days, one at the beginning of the project year (September 2012) and one towards the end (March 2013). Regular Skype calls involving the whole group took place about once every two weeks, and others involving the academic partners alone were also a regular feature. The outcome, as we shall see, was that a version of ‘patterns’ was developed which was significantly different from what has been produced previously. Notably, an initial decision was that ‘patterns’ were most likely to prove useful, at least in the first instance, if they were adapted from existing data in a circumscribed group of settings. The decision was made, reflecting on-going work by the design team, that the settings in question would be limited to those involving some form of (loosely defined) maintenance and construction engineering context.

### **4.3 Construction and Maintenance Engineering**

If some form of ‘generalised’ findings were to be an outcome of the work, it was important to gain an understanding of what the existing literature had to say and how problems in our chosen domain were typically framed. Our discussion of the literature is framed around some specific matters which have to do with the inevitability of failure and attempts to minimise its impact, the desire to enhance control so that failure rates are reduced or kept to a minimum and the recognised need to embed solutions within the actual work of maintenance. Problems, of course (and as all ethnographers can attest), can be seen to be at least partly characterised by the specific domains within which they unfold. This would seem to present a challenge regarding what lessons might be appropriately learnt from investigations in other areas of construction and maintenance. This does not preclude an examination of those features with a view to establishing what might prove to be relevant and valuable for design-related purposes. Our review took us through a variety of domains encompassing, inter alia, software maintenance, general ICT maintenance in education, manufacturing systems maintenance, software maintenance, industrial systems maintenance, aircraft maintenance, railway maintenance, ship maintenance, vehicle maintenance, road construction and maintenance, photocopier and printer repair maintenance, network equipment maintenance and home device assembly and maintenance work. Now, of course, such domains are heterogeneous. The purpose of this review of the literature was to provide an initial, high-level, view of common problems (if any could be found) and subsequently to identify a subset of domains which would constitute a manageable resource for comparative work. We were able to identify certain domain features which, if not universally shared, at least could be found in more than one setting. Thus and for instance, the literature on

software maintenance points to problems such as the need to keep systems running whilst maintenance takes place. As Swanson (1976) argues:

The amount of time spent by an organization on software maintenance places a constraint on the effort that may be put into new system development. Further, where programming resources are cut back due to economic pressures, new development is likely to suffer all the more since first priority must be given to keeping current systems ‘up and running’. (p. 492)

Swanson refers firstly to ‘corrective maintenance’, which mainly involves diagnosis procedures in relation to failure, and secondly to ‘adaptive maintenance’, which is done in anticipation of environmental change. Thirdly, there is ‘perfective maintenance’ which, in contrast, refers to performance enhancement (e.g., cost-effectiveness). The latter, he suggests, is not well understood but is at least as important as the first two since it is ‘directed toward keeping a programme up and running at less expense, or up and running so as to serve the needs of its users [and customers]’ (p. 493). Tan and Gable (1998) – also in the context of software engineering – demonstrate how there are radical differences between the attitudes of management and those of ‘maintainers’ and how this ‘knowledge gap’ also concerned what problems maintainers faced and how managers interacted with maintainers. Perhaps more importantly, from our point of view, there has been a developing recognition in the area of software maintenance that, along with attention to the above factors, we need to pay attention to local patterns of work. Bendifallah and Scacchi (1987) were amongst the first to recognise variation in local patterns of work. They argue (in a comparative study of two organisations) that we need:

... to understand the ways local circumstances in the workplace affect how and why people perform software maintenance tasks, and conversely, how maintenance work affects workplace arrangements. Local circumstances include the incentives and constraints for why people alter their software systems, and indicate when people act to maintain their systems. The workplace specifies where maintenance work is performed and the ways it is organized. How people order and perform their maintenance work also entails who does this work, and what kind of maintenance activity is performed. (p. 311)

Interestingly, the idea of ‘beacons’ or recurring patterns is occasionally referenced in this literature (see Boehm-Davis et al. 1996; Crosby et al. 2002). Ko et al. (2005) draw attention to the need for detailed studies of the way in which maintenance tasks are actually accomplished when Java programmers are using Eclipse, with a view to designing tools which might support their work. We might also briefly mention the work of Kemerer and Slaughter (1997) in this context, because they point explicitly to the gap between technical and managerial knowledge in maintenance work. They draw on the notion of patterns to argue that managerial information can be made more robust if these patterns can be identified. Hence:

Software maintenance is a task that is difficult to manage effectively. In part, this is because software managers have very little knowledge about the types of maintenance work that are likely to occur. If managers could forecast changes to software systems, they could more effectively plan, allocate workforce and manage change requests. But, the ability to forecast software modifications depends on whether there are predictable patterns in maintenance

work. We posit that there are patterns in maintenance work and that certain characteristics of software modules are associated with these patterns. (p. 1)

We then looked at maintenance in manufacturing industries. Bateman (1995) describes three basic forms of maintenance in relation to manufacturing. The first is *reactive*, such that equipment is allowed to run until it fails and then firefighting techniques are implemented. Gits (1992) calls this failure-based maintenance. The second is *preventive* (sometimes called use-based maintenance) and is obviously more proactive. Preventive maintenance usually involves maintenance work at particular and scheduled times. The third is *predictive maintenance* or is sometimes called condition-based maintenance. This relies on reporting about the condition of equipment and depends on new forms of diagnostic technique (measuring, for instance, vibration, noise, lubrication and corrosion). Laura Swanson (2001) introduces a fourth form, which she terms *aggressive maintenance*. She argues this is typical of strategies such as total productive maintenance (TPM) and which have been introduced as a result of global competition and enhanced technological capability. It is predicated on the view that small teams can create cooperative relationships by integrating production and maintenance functions. This, of course, implicates the raising of skill levels through constant communication and cooperation. Much of the literature is of a technical nature and focuses on such matters as optimization problems, replacement models and so on (examples are Nakagawa (2005) and Smith and Mobley (2008)). Other work focuses more on maintenance *management*. Here, the main interest is in best practice. Wireman (2004), for instance, notes that the failure to introduce proper policy and procedure is commonplace. He suggests that this is because maintenance is viewed by most organisations as a necessary evil and hence their main interest tends to be cost control. One of his observations is that, roughly, only half of time at work by maintenance workers is actually spent doing hands-on maintenance. He sees the main failures as having to do with job planning (better planning, he argues, leads to a cost ratio of 1:5 against poor planning) and with poor work order systems. Only 10 % of organisations, he maintains, have any form of performance monitoring or failure analysis. He further argues that much more maintenance work is reactive than it should be (evidenced by high levels of overtime). Frequently, the cause of these failures lies in under-skilling, lack of coordination between operations and maintenance, poor communication and poor use of materials (which are typically 40–60 % of all costs). One feature of this is overstocking and poor inventory control. Maintenance costs are also poorly assessed, since they seldom include the cost of ‘downtime’. Such emphases have led to a concern with ‘lean maintenance’ (see Smith and Hawkins 2004).

A third area where maintenance work has been theorised is in safety critical systems. Here, and unsurprisingly, measures of risk are paramount. The literature here is highly technical, although there are some studies based on managerial and ‘cultural’ perspectives (see, for instance, McDonald et al. 2000) where interviews and documentary analysis are used to understand how organisational cultures affect safety. The main thrust of their work is to show that professional subcultures mediate the rules and procedures mandated by organisations. The fact that failure

in safety critical systems can be catastrophic means that standard ‘trial and error’ approaches cannot be relied on. More importantly, there is a limited literature of a more sociological or social psychological kind which looks at the way in which groups in such contexts deal with problems. This includes some classic literature such as work by James Reason and Jens Rasmussen, work on ‘high reliability’ organisations by Laporte and Consolini (see, e.g., Laporte 1996; Laporte and Consolini, 1991, 1998a, b), work on aircraft cockpit errors by Charlotte Lind, more general theoretical work by Charles Perrow and – to a certain extent – work by the Lancaster school (Hughes et al. 1992; Randall et al. 1993) and others on air traffic control. Various, such studies show cultural factors of one kind or another produce ‘error tolerance’. Cultures, in this literature, either guarantee that errors are less likely to be made, are more likely to be seen or are less likely to be consequential. The exception is Perrow (1999) who argues that two factors – ‘interactive complexity’ and ‘tight coupling’ – mean that, in some contexts, accidents are more or less inevitable. ‘Tight coupling’ refers to the strong causal link between actions or events, whilst ‘interactive complexity’ refers, fairly obviously, to the many and varied ways in which people and systems can interact in complex organisations. Perrow is particularly interesting insofar as, when he points to failure, he identifies several individual factors which are relevant, including human error, mechanical failure, the environment, design of the system and the procedures used. Perrow is nevertheless insistent that catastrophic failure is never produced by any one of these factors but by the interaction of them all.

What became clear from our review of these and other domains was that there are some fundamental perspectival differences in play. Maintenance problems can be seen in a variety of ways and one of them we might call ‘managerial/technical’. Maintenance has increasingly been seen through the organisational lens and associated with this as a management issue. Maintenance management, in other words, has become a quite distinctive approach, characterised as:

All the activities of the management that determine the maintenance objectives or priorities (defined as targets assigned and accepted by the management and maintenance department), strategies (defined as a management method in order to achieve maintenance objectives), and responsibilities and implement them by means such as maintenance planning, maintenance control and supervision, and several improving methods including economical aspects in the organization. (Marquez 2007, p. 3)

The above demonstrates a view of maintenance which is, in the main, top-down. In other words, even though domains vary, the general perspective is that defined goals, along with procedures which are consistently implemented by managerial teams, define ‘best practice’. There is, even so, an extensive literature which argues that best practice is not easily achieved. Various barriers need to be overcome (see, e.g., Raouf and Ben-Daya 1995). Cooke (2000) argues that certain organisational barriers are consistently found and they include political, financial, departmental and inter-occupational barriers.

In recent years, much of this work has been conceptualised in terms of trust and dependability. Avizienis et al. (2001) characterise dependability in the following

way (much of what they say relates to computer systems, but there are some overlaps):

*“A systematic exposition of the concepts of dependability consists of three parts: the **threats** to, the **attributes** of, and the **means** by which dependability is attained”*, and argues that fault prevention, tolerance, removal and forecasting are the critical issues. The threats to dependability, unremarkably, consist in faults, errors and failures. Nakagawa (2005) also emphasises reliability in maintenance and argues that the basic attributes required are availability (readiness of service), reliability (continuity of service), safety, confidentiality, integrity (absence of improper system state alterations) and maintainability.

In contrast, there is an approach we can think of as broadly ethnographic. An acknowledged classic in this area is Julian Orr’s (1996) study of photocopier repair. Orr makes a number of related points, many of which are still under-examined by the more ‘managerialist’ literature. They include the fact that *work is heterogeneous*, meaning that problems are much more varied and unpredictable than we might assume. He further identifies, along with other ethnographers, the *social distribution of expertise* (he does not call it this; see Randall et al. 2007) meaning that we cannot assume that all are equally skilled. He suggests that not all maintenance work is done in the same way. It can be done in order to provide a minimum level of serviceability – i.e., doing the bare minimum – or it can be aimed at actually providing a solution to problems. He makes the point that decisions of this kind are made because of a variety of pressures including time available. Probably most importantly, he identifies the fact that this work is done collaboratively, even if maintenance is being done by one person at any given moment. They do this in large part by informal knowledge sharing, for instance, by meeting up for breakfast. Orr makes the point that much of this talk is generated by the sheer unpredictability of their work. Each machine behaves differently. Problems as described sometimes turn out not to be the problem at all. What is important here is the fact that much of the knowledge and skill is not formally codified but depends on experience and practice. Technicians become familiar over time with the characteristics and weaknesses of particular models of machinery. Technicians in this field have values concerning ‘being good at your job’ which include thoughtfulness, attention to detail, freedom from panic and resourcefulness when the documentation does not provide the answers (p. 34). Orr also reflects on the codification of expertise, pointing out that such codification is seldom entirely satisfactory. It seems that, in this context, documentation did not point adequately to solutions. In a similar vein, Carstensen (1999) has pointed to difficulties inherent in codifying processes when perspectives, practices and terminology are different from work group to work group. For a critique of the whole approach, see Schmidt (2012). Schmidt uses an example from the manufacture of diesel engines for ships:

... diesel engines for large ships are not produced for inventory, in the hope of somebody filing an order sometime in the future: marine propulsion plants are far too costly and take up far too much space for that to be a viable business model. Thus, many months may pass between orders for an engine of any particular model are received. Consequently, it was likely to be difficult for whoever was tasked with a given assembly job to remember how exactly to assemble the particular model requested in the given purchase order. So, although at least some of the workers would know how to assemble the model in question,



each of them would not necessarily be able to assemble it efficiently and correctly. They therefore strove to maintain their collective capabilities by documenting the sequential order and key operations of the procedure and by thus offering means for each other (or for future colleagues) to maintain (or acquire) their individual ability to do so efficiently and correctly.

Schmidt's work, *inter alia*, draws attention to timeliness, to 'competent readership' and to practical matters of one kind or another. Pipek and Wulf (2003) provide a critical review of the problem of knowledge sharing in maintenance work in relation to what is sometimes called 'organisational memory'. This describes the problem of retaining the expertise and experiences of skilled members. Pipek and Wulf draw on the work of Mark Ackermann who identified a number of issues in relation to the informalities of work and the difficulty of knowledge sharing. These variables include the complexity of the knowledge domain, the interactive nature of the problem-solving process, technological infrastructure, the existence or not of a common body of knowledge (shared expertise) and how dynamic changes in the body of knowledge are. Pipek and Wulf describe difficulties in maintenance work in a steel mill which have some similarity to those described by Orr, mainly in relation to documentation of various kinds and how accurately it describes the real state of play. They point to some typical issues around the nature of the documents (paper and electronic), which include that a large number of documents are not classified or are classified in such a way they cannot easily be found, that many are old or of poor quality, that search functions are inadequate, that some knowledge resources are 'private' and so on.

None of this is to argue that codification is impossible, only difficult. Indeed, work at Xerox following on from Julian Orr's study involved precisely the design of knowledge-sharing systems for photocopier repair across international sites. The work on this (see Bobrow and Whalen 2002; Yamauchi et al. 2003) is interesting because it discusses the organisational barriers to getting ethnographic work accepted and also provides examples of cost and productivity benefits that resulted.

Our review of the literature indicated a number of things to us. Firstly, managerialist and technical approaches, whilst pointing to the need for actions at the strategic level, for instance, in relation to the 'the acquisition of the requisite skills and technologies' or to 'the correct assignment of maintenance resources (skills, materials, test equipment, etc. to fulfil the maintenance plan). (Marquez 2007) say little about how such things might be achieved in practice. Similarly, technical approaches focus in the main on metrics and tools for the measurement of known criteria. Both, we feel, produce generalised levels of argumentation that are insufficient for dealing with the specific requirements identifiable in any given setting. The 'ethnographic' approach, and readers would be familiar with this argument, tends instead to emphasise local contingency, the gap between formal rules and procedures and actual practice and so on. A significant 'background' assumption here is that the kinds of problem – as well as the practical management of solutions – that occur are specific to the setting under investigation and are best understood by close investigation of the cooperative elements of work. However, and as we have already indicated, the practical management of ethnographic enquiry

is problematic. A number of different individuals, remember, may be working in a number of different settings and – sometimes at least – taking a considerable period of time to ‘gear in’ to those settings (which are frequently of a highly technical nature). They in turn are reporting back to a large and disparate group of designers who sometimes struggle to understand the results. The task the group set itself, then, was firstly to find a means to provide a manageable level of generalisation, one which both indicated what kinds of commonality might exist across a range of broadly similar settings but at the same time retained a sensitivity to local contingency. Secondly, an aim was to find a language which allowed for a common means of expression, such that both ethnographers and designers might find it easier to interrogate their own and others’ work.

## 4.4 Patterns

The idea of the pattern language originates with Christopher Alexander (1979) but was quickly taken up within software engineering communities. Alexander was offering an approach to the design of buildings which reflected patterns of use and as such was critical of ‘heroic’, modernist conceptions of architecture which rather ignored user needs in favour of abstract notions of ‘function’. In ‘a pattern language’, Alexander claimed to have identified 253 basic patterns which recur in architecture. Patterns have particular attributes, which he characterised as follows:

Name: A name to identify the pattern.

Context: The situation(s) where the pattern is relevant.

Forces (problems): The forces present which may constrain or suggest alternative solutions. When these forces are in tension with one another, the problem is harder to solve and a compromise may be necessary.

Solution: A solution which resolves, as far as possible, the various forces.

For Alexander, these patterns had to be validated in some way (including observation) and were prescriptive. That is, they instruct us what to do about certain things. We should not forget that these patterns were also invariant. In other words, they are to be found across settings (see Dearden and Finlay 2006, who also provide a useful summary of the way patterns are used in a variety of contexts). They can be distinguished from other forms of generalisation, they argue, in the following ways:

1. the level of abstraction at which guidance is offered;
2. the grounding of patterns in existing design examples, or “capture of practice”;
3. the statement of the problem addressed by a pattern;
4. the discussion of the context in which a pattern should be applied;
5. the provision of a supporting rationale for the pattern;
6. the organisation of patterns into pattern languages; and
7. the embedding of ethics or values in the selection and organisation of patterns. (Dearden and Finlay 2006, p. 21)

There has been a significant uptake of ‘patterns’ in computer science (see, e.g., Gamma et al. 1993; Gabriel 1996; Rising 1998; Denef 2012), but here the idea of ‘pattern language’ became much more significant. According to Erickson (2000),

patterns are particularly interesting for object-oriented programming because they embody concrete prototypes, are grounded in the social and express values. Patterns of code are easy to identify, but connecting them into more complex frameworks requires a language which consists of connective rules. Salingaros (2000) gives some examples:

- One pattern contains or generalizes another smaller-scale pattern.
- Two patterns are complementary and one needs the other for completeness.
- Two patterns solve different problems that overlap and coexist on the same level.
- Two patterns solve the same problem in alternative, equally valid ways.
- Distinct patterns share a similar structure, thus implying a higher-level connection.

Patterns, then, have a formalising quality, involve connections and levels and further involve selections as to what is and is not relevant. It has been pointed out that one of the main values of the pattern approach is that it documents or codifies these relevant factors. Put simply, they have a systematic quality. If patterns can be systematically connected through the use of a common language, then, they should in principle provide us with a framework to anchor design. Nevertheless, for the reasons given above, they cannot determine design. There is an obvious analogy here between software engineering approaches which rely on formal models of organisational ‘need’ and those which orient to user behaviours, since understanding patterns clearly entails understanding behaviour (as well as functionality) in a systematic way. Having said that, the idea of a ‘pattern’ is a rather vague one and can mean many things. The notion of ‘patterns’ has also been subjected to extensive critique (see Dovey 1990) insofar as it is clear that patterns cannot provide complete design solutions (the reasons for this have to do with their putative inability to encompass economic, policy and construction or implementation issues). They can, in other words, only be seen as one tool in the toolbox. One of the evident difficulties in the application of patterns to human behaviour lies in the different perspectives on ‘relevance’ that are brought to the table. By this, we mean that patterns are sometimes applied within disciplinary boundaries but might be better understood as a communication device for interdisciplinary work. Deneff (ibid), for instance, deploys patterns in an attempt to integrate the perspectives of ethnographers and designers in relation to firefighting. Mahemoff and Johnston focus on patterns of *usability*. Their patterns are structured at the level of task, user, user interface and ‘entire systems’. In a different vein, and of more immediate relevance for our purposes, Martin and Sommerville draw extensively on existing ethnomethodological work to draw attention to recurrent topics. These, for ethnomethodologists, might include:

1. Sequentiality and temporality
2. The working division of labour (egological-alteriological principles)
3. Plans and procedures (representations)
4. Routines, rhythms and patterns (orderliness in self-organising systems)
5. (Distributed) coordination
6. Awareness of work
7. Ecology and affordances

What Martin and Sommerville try to do is use these ‘topics’ as a resource for generating patterns in more specific circumstances and then exemplifying the patterns through data. They argue:

Patterns are intended to be a resource that is a structured collection of findings from field studies of work and technology. As such, reading through them should provide a good background understanding of some of the social design issues that arise out of these ethnomethodologically-informed studies.

Now, Martin and Sommerville accept the Alexandrian principle that patterns are independent of context and can be deployed across a range of different settings. Our position evolved differently. Partly as a result of the work that had already been done in various settings by the design team and partly because of the similarities and differences we glimpsed during the literature review process, we were concerned to circumscribe our generation of patterns by limiting them to settings where we were confident we had sufficient knowledge and which were, on some level, similar.

Equally important is the discussion of the way in which patterns can be integrated into approaches to observation, intervention, participation and learning, and communication. There is a certain amount of work on the relationship between pattern languages and participatory design (PD) (this includes, for instance, Borchers (2001); Lin and Landay (2002)) and for interaction design (see, e.g., van Welie and van der Veer 2003). Other work, like that of Goodyear (2005), orients to education. For our purposes, this is a critical feature. It was and is a fundamental feature of our enquiries that they should aid the communication process – ethnographer to ethnographer and ethnographer to designer (and vice versa). For our purposes, there were nevertheless decisions to be made.

## 4.5 Developing ‘Patterns of Work’

Our strategy was to identify some high-level, domain-independent questions which could then be decomposed and translated into a domain-relevant set of questions. We make no claim to originality here. The questions are similar to those in Martin and Sommerville (2004), to those posed by Randall et al. (2007), to Checkland’s CATWOE analysis and so on. They were evolved by the academic partners and were delivered with some rough sub-questions which were intended to be illustrative of the range of issues. It is important to recognise that these high-level questions have no theoretical status – they are simply orienting devices. They were in no sense intended to direct enquiry but were rather geared to the assumption that the evolution of our patterns would be an iterative process and one which ethnographic data would feed into. That is, ethnographic findings and evolving patterns would be mutually elaborated. Indeed, they were posed as questions precisely for this reason. The ‘nine questions’, as they became referred to, are laid out below, exactly as they were posed to the project group.

1. *What actors are involved?* Of course this raises a number of important sub-questions, especially in the light of the fact that work is increasingly complex, mobile and Internet reliant. Do we include only immediate workers in our studies, or should we be looking at a wider organisational context? When knowledge is shared, is it only shared locally? How heterogeneous are the groups that share a working environment?
2. *What are they doing?* Again, we need to make decisions about granularity, and this depends on how we decide on what is interesting. For some purposes, it is obvious that we need to pay attention to the sequentiality of work in some detail (in what order do people actually do things?). Video analysis is an obvious way of doing this. For other purposes, less detailed observations may be possible, and for others interviews are likely to elicit information that observation will not ('Why did you do that? Why not do it another way?' 'What was the problem there?').
3. *Why are they doing it?* This speaks to the goals of management and the organisation, to 'accountability' issues and so on.
4. *Why are they doing it this way?* In much the same way, there may in principle be any number of different ways of performing tasks. If things are typically done in one fashion, then we can ask why. Are there constraints caused by material resources? How rule governed are the activities? Are there problems with skill levels? Is this the most elegant and effective way things can be done given various constraints?
5. *What materials, resources and spaces do they deploy and in what order?* The point of this is obvious, but we need to remember that ALL material resources are relevant. Scraps of paper like Post-it notes, scribbled bits of information, reminders and other more human resources (e.g., asking questions) are important.
6. *What knowledge and skill do they demonstrate?* This is possibly the most difficult thing to uncover, not least because it often means the ethnographer has to understand technical terminologies. It also means one has to pay attention to a series of quite mundane skills – who seems to know most and what is it that they know (sometimes called the social distribution of expertise).
7. *What (if anything) needs to be changed?* It is often easy to identify bottlenecks, problems and so on. It is not quite so easy to see how to correct them. An example is the idea of redundancy. In one sense, if processes are redundant, it means they are being duplicated for no good reason. It does, however, mean that we have to be sure there is no good reason. It could be, for instance, that redundancy is how mistakes are identified and corrected.
8. *How do we go about justifying and making changes?* Whatever problems we identify, there is still the important issue of persuasion. Management is often reluctant to accept that the themes the ethnographer has identified are real (I have personal experience of this). Workers sometimes have a vested interest in preserving current work patterns.
9. *How do we evaluate the effects of intervention and over what period of time?* One of the classic ethnographic problems is that of evaluation of change. Over what period of time should we be looking at practices and the changes that are taking place?

Following literature reviews and the generation of the ‘nine questions’, the group met at a workshop to interrogate existing data with a view to establishing what general areas of interest in the field sites might be identified and hence what the focus of subsequent patterns might be. It was agreed that the project, which was to last 2 years, would be divided into two fairly distinct phases, and only the first six questions above would be dealt with in year 1. That is, the first phase of the project was regarded as largely descriptive/analytic, whilst the second phase was to be representational/interventionist. As pointed out above, an early decision had to do with the practical possibilities inherent in making comparisons from field research conducted by the team allied to published literature. It was felt that too much heterogeneity might result in generalisations that were of little practical use, and so settings which could loosely be termed as ‘engineering maintenance’ or ‘construction’ were chosen. Following this and again done by pulling out some identifiable features from data we already had (based both on ethnographic studies which were under way or, in one case, completed and on the literature we had examined), we cautiously isolated five themes which formed the basis for the development of the patterns. During this workshop, members of the group described the work they were doing in various settings, including train maintenance, power plant construction, an Australian mining camp and an international software collaboration. From this we tentatively outlined five potential patterns:

- (a) Finding tools and materials
- (b) Sequencing technical activities
- (c) Coupling work activities
- (d) Sharing knowledge
- (e) Scheduling for contingencies

These general patterns were not clearly demarcated, at least to begin with, but reflected a common-sense approach to the kinds of issue that we saw as typically arising in the settings under investigation. ‘Finding tools and materials’ is fairly self-explanatory and refers to the fact that tools and materials are not always easily identified, can go missing and are sometimes used by more than one individual or group and that, as a consequence, work can sometimes be held up. ‘Sequencing technical activities’ reflects a classic ethnomethodological concern, the detailed description of how activities are organised on the basis of an egological orientation. It deals, that is, with the ‘what do I do next’ questions rehearsed in Randall et al. (2007). ‘Coupling work activities’ deals with the fact of coordination ‘down the line’. That is, how the activities of one group of workers cascade down consequentially to those of another. ‘Sharing knowledge’ was an attempt to specify the skills, expertises, local knowledges and so on that might be possessed by one group but not necessarily shared with others with a view to understanding how consequential that might be. Finally, ‘scheduling for contingencies’ reflected the well-known fact that the standard ordering of work schedules was sometimes disrupted by contingencies, by new priorities and so on.

It needs to be stressed here that data collection from various sites was continuing. Additional data was analysed and fed into the evolution of the patterns. Space

precludes a detailed analysis of the data from the several studies that fed into our decision-making, and in any case the provision of detail is not the primary function of the chapter. We can, however, at least give some examples of the way in which ethnographic data fed into the refinement of our comparative patterns. Below, we give an outline of two. Again, we need to be clear that the purpose of this chapter is not to rehearse ethnographic data in any detail, and we do not do so. Our purpose is to do with representing ethnographic data in such a way that it can be seen to ‘fit’ with evolving schema.

*Example 1* One of the areas that Hitachi fieldworkers had been studying was train maintenance in the UK. Maintenance work of this kind is done in a depot and, at the site in question, on five tracks in the depot. Five teams of operatives work in parallel, each consisting of about 15 maintainers. Most of the work done on site (though not all) is routine and scheduled, and, typically, a manual of rules and procedures is used to identify the stages of specific operations, and work completed is recorded on an application running on a laptop. In presenting the work, fieldworkers drew attention to the fact that work was frequently held up as a result of the fact that various tools and materials were not always ‘to hand’. There were a number of different reasons for this, including the fact that manuals did not provide exhaustive lists of tools required for the completion of particular jobs which meant pauses whilst the requisite tools were obtained, the fact that the stock management system was not easy to use and inventory checking was cumbersome and that tools which were supposed to be located on workbenches often went missing. Now, in comparing this site to others, we observed that ‘finding tools and materials’ was a generic problem, though one which took on specific characteristics in different circumstances. Thus and for instance, a range of fairly typical dimensions seemed to be implicated. These included such issues as whether the same people who organised resources also used them, whether adequate catalogues of resources were maintained, where equipment was kept, whether more than one person or team used the equipment and so on. After a process of iteration, the following pattern was evolved, with the issues that seemed to be salient highlighted:

Now, what we describe here is not a pattern in the sense that it was used by Alexander. It is not prescriptive and it does not seek to describe universalities. Indeed, the questions are designed to elicit the dimensions of variation. The pattern itself, however, is not the point. We should remind ourselves that the purpose of the ‘patterning’ process is to make complex data from individual sites available to both other ethnographers and to designers in a form which enables them to ‘read’ the data in a usable form. To this end and having said that data feeds into pattern construction, ethnographic data is also progressively being structured in a way which is consistent with the patterns. The process, in other words and as already stated, is mutually elaborative. An example is given below (Tables 4.1 and 4.2).

We should remember here that it is the ethnographic data that provides for the specific rendering of the problems experienced at work in this context as well as, of course, descriptions and analyses of the ordinary routines. The patterns do not, and cannot, replace the insights the ethnography provides. The pattern, that is, represents

**Table 4.1** Maintenance engineering: finding tools and materials

The questions	<i>The pattern</i> -maintenance work-finding tools and materials	Comments
What actors are involved?	1.1 How are maintenance teams constructed?	1.2.1 Do they deal with planned work or exceptions?
	1.2 What is the division of labour?	1.2.2 How are tasks allocated?
	1.3 How is authority formally distributed?	
	1.4 How are responsibilities actually organised?	
	1.5 What aspects of the work involve fixed positions and/or mobility?	1.5.1 Do they work in confined spaces? 1.5.2 Do they need to cover large distances?
2. What are they doing?	2.1 What different kinds of equipment are used?	2.1a How specialised is the equipment and does it take specialised knowledge to use it?
		2.1b Are the same people responsible for the organisation of resources as the people who use them?
		2.2.1 Is equipment kept in a centralised position?
		2.2.2 Is a catalogue of the equipment used?
		2.2.2a Is the catalogue easy to use?
		2.2.3 Who maintains the catalogue and is it done accurately?
		2.2.4 Is the catalogue digital or physical in form?
		2.2.5 Who allocates equipment?
		2.2.6 Who is responsible for its return?
2.2 How is equipment organised?		2.2.7 How often does equipment go missing or get lost?
		2.2.8 Is equipment used by others?
		2.2.9 How long does searching for equipment take?



<p>Why are they doing it? Why are they doing it this way?</p>	<p>3.1 How adequate is the equipment they use? 3.2 How clear are instructions for using it? 3.3 Is there anyone with specialised skills to use the equipment? 4.1a Are there health and safety concerns? 4.1b Are there external reasons for organising this way? (e.g., do contracts specify what equipment or information resources will be used?) 5.1a Are equipment and materials standardised?</p>	<p>3.1.1 Is equipment up to date? If not, why not? 3.2.1 Do people have difficulty working with the equipment? 3.3.1 Are the members of the team adequately skilled for equipment use? 4.1.1 What restrictions are there on the way equipment can be used? 4.1.2 Are these restrictions imposed by the government, the organisation or trade unions? 5.1.1 Do workers carry personal tools or materials with them? 5.1.2 Do workers carry personal information with them? 5.1.3 Are documents in standard formats? 5.1.4 Is documentation in physical or digital form? Why? 5.1.5 Does it slow the work up? 6.1.1 Are the specialists on site? 6.1.2 Are there enough of them? 6.1.3 Do they have adequate knowledge of local conditions?</p>
<p>What materials, resources and spaces do they deploy?</p>	<p>5.1b Do workers ever use 'unofficial' equipment? 5.1c Is the equipment always available? If not, why? 5.1d Are information resources/documents used? 5.1e Are there any overheads to using this equipment or information?</p>	
<p>What knowledge and skills do they demonstrate?</p>	<p>6.1a What kind of training is needed to use the equipment? 6.1b Are there specialists?</p>	

**Table 4.2** Finding tools and materials: train maintenance

Case 1	Source	UK train maintenance
	1.2.3 How are tasks allocated?	Maintenance teams are mainly concerned with scheduled, routine, maintenance, although they have to deal with breakdowns as well.
	2.1b Are the same people responsible for the organization of resources as the people who use them?	There is a central storage area, but tools are only returned once a day. There are other storage areas located at different places in the shed. There is no clear responsibility as to who should return tools to the designated areas.
	2.2.1 Is equipment kept in a centralised position?	Equipment is kept in a variety of places and it is sometimes time-consuming to walk around trying to find tools. Necessary tools/ materials are often missing in the storage area. Expendables like cans are sometimes placed in the office or work bench, not in the designated storage area.
	2.2.2 Is a catalogue of the equipment used?	Yes
	2.2.2.a Is the catalogue easy to use?	It is difficult to use. Workers do not use the same language to describe tools and equipment as is used in the e-catalogue. It can be very difficult to trace what new equipment is needed.
	2.2.4 Is the catalogue digital or physical in form?	Digital
	2.2.5 and 2.2.6 Who allocates equipment and who is responsible for its return?	The store manager at the beginning of the day, otherwise, individual workers as needed. No-one has a specific responsibility for return
	2.2.7 How often does equipment go missing or get lost?	Frequently
	2.2.8 Is equipment used by others?	Yes, all the time. Teams are working in parallel and often borrow equipment from other teams. Some equipment is in short supply
	2.2.9 How long does searching for equipment take?	It depends, but we have seen examples of 2 h

a means to represent problems in a particular way, using a consistent language and thus affording an easy and convenient means to compare.

*Example 2* A second theme, which we evolved by drawing on existing data, was that of ‘coupling work activities’. This was intended to describe situations where the work of one individual or team affects the work of another individual or team. The degree of ‘coupling’ of work activities is scarcely a new insight (see, e.g., Weick 1976; Orton and Weick 1990), but our purpose, we remind ourselves, is pragmatic. That is, we set out to establish what the conditions which in practice affect the flow of work from one group to another might be. Again, we drew on

data from published (and some unpublished) work conducted by others and also on fieldwork data collected by the design team. The questions that we evolved, again iteratively, as new data became available, are listed in Table 4.3. An example of the data which informed the construction of the pattern comes from analysis of construction and maintenance work in a power plant. Here, members of a warehouse management group make preparations in response to requests for receiving (from a subcontractor) and delivering (to the construction site) materials (two people for incoming materials and three in charge of delivery). The materials in question include pipe lengths of varying and non-standard sizes and shapes which need to be stored in a physical location. The amount of space available for the storage of these pipes is limited. The pipes are ordered from a subcontractor some 3 months in advance of need. Having said that, changes in construction schedules mean that pipes are often stored for longer than that. The locations at which pipes are stored are registered using a GPS system. Problems occur, however, when the warehouse team unexpectedly receives materials from the subcontractor which do not correspond to digital delivery slips. They have to be inspected against load manifests, recorded and then stored somewhere. Having said that and to compound the problem, the pipes that are delivered will not always correspond accurately with descriptions on the manifest. That is, they will often not be needed at that moment because work schedules have been changed. Thus, they need to be stored, often for months. The pipes are too big to be handled entirely by hand and a range of moving equipment is used. Moreover, because they are often of unusual shapes and sizes, they sometimes cannot be placed in spaces originally allocated for them. There is a limited amount of space on site for the storage of these pipes and, because they are often stored for months on end, there is a tendency for them to be moved around as workers search for pipes that are needed in the near future. Limited time means that the location of pipes is not always accurately recorded.

In theory, pipe delivery to the construction site is organised three days before fitting specialists pick them up. At the same time, pipe fitters sometimes make unexpected demands because of changes to their work schedules. These sudden and unexpected requests mean that preparation is sometimes hurried and, more consequentially, that pipe fitters from the construction site collect pipes themselves (often displacing other pipes whilst they search for what they need, making it difficult for the warehouse group to control the location of inventory). Here, then, the practices of one group (pipe fitters, suppliers) have a significant effect on the efficiency of another (warehouse management) (Figs. 4.1 and 4.2).

Here, then, ethnographic data can be described in relation to the highlighted features in the pattern. That is, the data is examined to see what the salient questions might be, as below:

In a similar vein, then, the problems that could be identified across the different sites involved in this work (most of which are not reported here) are revealed by detailed ethnographic analysis. The pattern, once again, is used as a device for producing results in a format that allows for comparison.

**Table 4.3** Coupling work activities: pipe delivery and collection

The questions	<i>The pattern</i> -maintenance work-coupling work activities	Comments
1. What actors are involved?	1.1 How are maintenance teams constructed?	1.1. 1 Here, we want to know about the numbers of people in teams
	1.2 What is the division of labour?	1.2. 1 What jobs are people formally required to do? What are their job descriptions? Do they actually do more than that?
	1.3 How is authority formally distributed? 1.4 How are responsibilities actually organised?	1.3.1 and 1.4.1 The kind of thing we are interested in here is whether members of the team work strictly to orders or whether they themselves have discretion in what they do. Can they, for instance, change the order of work if they think it is efficient to do so without asking anyone else?
	1.5 What aspects of the work involve fixed positions and/or mobility?	1.5.1 Do they work in confined spaces?
		1.5.2 The literature indicates that geographical mobility creates a different set of problems than working in fixed locations, and we have found the same thing. Some problems can be caused by having only a small amount of space to work in, others by having to work a long way from other people
2. What are they doing?	2.1 What order is the work done in?	2.1.1 Here we are specifically interested in sequentiality-detailed description of the way work is done and if possible the reasons why people decide to do it like this. Video and pictures are often used to get a good sense of this
		2.1.2 Is there a difference between expected workflow and actual flow of work?
	2.2 How rule based is it?	2.2.1 How much discretion is there in the way people do their work, or are they tightly controlled?
	2.3 How coupled with other work is it?	2.3.1 Do changes cascade through to other work? Do stoppages in one place cause stoppages in another?
	2.4 How is it coordinated?	2.4.1 How are problems communicated to other work groups, if at all? How are changes in work plans communicated, if at all?

(continued)

**Table 4.3** (continued)

The questions	<i>The pattern</i> -maintenance work-coupling work activities	Comments
	2.5 What is difficult?	2.5.1 What things actually cause plans to be altered? 2.5.2 How often do they occur?
	2.6 What is unexpected?	2.6.1 We make a difference between ‘normal natural troubles’, the things which regularly occur which affect the flow of work, and the ‘unexpected’, more serious, perhaps more rare, events
3. Why are they doing it?	3.1 What does the organisation demand of them (workload)?	3.1.1 What time and resource allocations are made?
	3.2 How clear are the instructions?	3.2.1 What is the role of manuals and other instructions? Do workers pay attention to them? It is well known that workers often find it difficult to identify what the problem is in a manual (or where to find the answers) and often find other solutions
	3.3 Who instructs them?	3.3.1 What is the structure of authority and who, in practice, decides what should be done next?
	3.4 How are roles actually distributed?	3.4.1 Is there any difference between formally described roles and what people actually do?
	3.5 How formal/informal is the distribution of tasks?	3.5 Do managers or supervisors change the order of work, or do they stick rigidly to planned schedules?
4. Why are they doing it this way?	4.1 What determines the order they do things in?	4.1.1 Why are there sometimes delays?
	4.2 How easy is it to organise and coordinate?	4.2.1 What don't they know that they need to know?
	4.3 How well does resource management fit the flow of work?	4.3.1 Are there resource problems because no one in management knows what is needed?
	4.4 How economic is it? Is there waste or redundancy?	4.4.1 Can we tell the difference between unnecessary duplication and useful checking?
	4.5 What disrupts the work, how long for and who is affected?	4.5.1 What are the factors that create ‘normal natural troubles’?
	4.6 What external factors govern the work?	4.6.1 Work activities are often constrained by factors such as contract obligations, health and safety regulations, legislation and other background conditions and external scheduling

(continued)

**Table 4.3** (continued)

The questions	<i>The pattern-maintenance work-coupling work activities</i>	Comments
5. What materials, resources and spaces do they deploy?	5.1 What limits the pace of work?	5.1.1 Here, we are mainly thinking of technical limitations – are the tools they use adequate?
6. What knowledge and skills do they demonstrate?	6.1 What do workers need to know and how are they trained?	6.1.1 To what extent is the training on the job?
		6.1.2 Is it regularly updated? How is skill maintained?
	6.2 When are these skills obvious?	6.2.1 When can you tell that workers need high levels of skill and when does it matter less?
	6.3 How technical is the work?	6.3.1 How difficult is it for the ethnographer to understand what the skills are?
		6.3.2 Are the skills social as well as technical?
	6.4 How easy is the skill to acquire?	6.4.1 How much training is needed?
		6.4.2 Do people pick up the skills by informal methods?
	6.5 How much of it needs to be shared?	6.5.1 Is it enough that there is one person who has the skill and can tell others what to do, or is it important that everyone has it?
6.6 What is the distribution of skill?	6.6.1 Is everyone equally good at the work, or are some people obviously more expert than others?	
6.7 How is information maintained and by whom?	6.7.1 Does a knowledge base need to be maintained, and if so, how is it done?	
1.5.1 and 1.5.2	Storage space available is limited. The problem is made more difficult by the fact that pipes are of many different shapes and sizes. These are often non-standard which means they cannot easily be stored in a defined order	
2.3.1 Do changes cascade through to other work? Do stoppages in one place cause stoppages in another?	The problems cascade across three locations. They start in the subcontractor's, where pipes are supposed to be made in accordance with work schedules but unexpected, and urgent, requests are quite common. In turn, this means that pipes are delivered to the storage site in a somewhat unpredictable way and sometimes have to be stored for long periods. Their non-standard shape and size means they are difficult to organise and other pipes have to be moved to accommodate them	
	When pipes are to be delivered to the construction site, they are often difficult to find because they have been moved, which causes delays in construction	

(continued)

**Table 4.3** (continued)

The questions	<i>The pattern-maintenance work-coupling work activities</i>	Comments
2.4.1 How are problems communicated to other work groups, if at all? How are changes in work plans	There is very little communication across different workgroups, except through documentation. Changes in storage location as pipes get moved around are not communicated at all	
2.5.1 What things actually cause plans to be altered?	Changes to work schedules at the construction site cause most problems  The contingencies which arise on the construction site mean that the order in which the pipes are manufactured is not the same as the order in which they are needed, meaning that pipes have to be stored often for long periods	
2.5.2 How often do they occur?	Frequently. It is quite common for work to be delayed because the right pipes or materials cannot be found	
2.6.1 We make a difference between 'normal natural troubles', the things which regularly occur which affect the flow of work, and the 'unexpected', more serious, perhaps more rare, events	Most of what happens can be described as normal, natural troubles	
4.1.1 Why are there sometimes delays?	The delays are almost entirely caused by the different problems associated with storage	
4.2.1 What don't they know that they need to know?	There is very little communication between one group and the next, other than through documentation	
4.3.1 Are there resource problems because no one in management knows what is needed?	The biggest resource problem is lack of space at the site	
4.6.1 Work activities are often constrained by factors such as contract obligations, health and safety regulations, legislation and other background conditions and external scheduling	The various delays are caused by differences between schedules, which are established months in advance, and the immediate needs of the construction site	

## 4.6 Conclusion

The work we report above was conducted, as indicated, with some pragmatic ends in view. It was intended to provide, as we have intimated, a common set of questions which could be decomposed to address specific themes relevant to the setting in question but phrased in a way which was generic enough that comparisons with other settings could be made. It was done with the approximate aim of achieving

**Fig. 4.1** Storing and retrieving pipes



**Fig. 4.2** Storing and retrieving pipes

certain quite pragmatic objectives. These included shortening the time taken to do ethnographic work by sensitising fieldworkers to ‘what they might find’; addressing the concerns of designers who had difficulty understanding what the lessons of fieldwork data might be; and perhaps most importantly providing a focus on problems of similarity and difference. Design work in the plant construction case, for instance, is further forward than in the train maintenance case and has focused initially on using an iPod touch for video conferencing and for the sharing of visual images. Visual imagery, in the context of non-standard sizes, seems to be a great deal more effective than any other form of description. Medium term design is orienting towards an augmented reality system which will dovetail with a system intended for use in construction sites themselves and which will feature point and tag functionality such that visual images of locations are overlaid with other data. The point here, of course, is that it would be naïve to imagine that envisaged systems will be used in one setting only. As far as possible, for sound economic reasons, they will be deployed in a range of settings where similar issues are described. The patterns, as we have remarked, are intended to outline the lines along which similarity and difference can be identified. Whilst we are not the first to remark



on the tensions between the provision of detail inherent in case studies and the comparative work, we do believe that few efforts have been made to examine these issues *across cases*. What we have described above is an on-going attempt to do so. A significant part of the work has been concerned with finding cases at an appropriate level of generality, such that useful comparison can be made. Although the conceptual issue of what exactly we might mean when we talk about ‘domains’ and ‘settings’ was not a part of our deliberations, the selection of cases which were ‘similar enough’ clearly was a relevant and problematic matter. We continue to examine the patterns in the light of data from new cases. A recent effort has been the examination of a railway control room.

Having said this, the patterns undergo constant refinement. The design team finds them useful but, at the same time, expresses certain reservations. To some extent, this is because we have not always had a clear, shared, understanding of exactly what benefits might accrue from the work. For instance, it became progressively more clear over the year that one implied (but initially unexpressed) need was to enable inexperienced workers to go into the field armed with something more than an ‘ethnography and how to do it’ literature, something which the patterns were never intended to do. Equally, the patterns were perceived to overlap such that it was sometimes difficult to identify which pattern asks pertinent questions about which situation. Even so, fieldworkers involved in the business of representing their work to designers report that they feel more confident in their efforts to do so. As one of them said, ‘at last, I feel I have a language I can use to them’. Regardless, how best to represent these evolving structures such that both requisite detail and necessary generality are encompassed remains an issue. Representing those similarities and differences and aligning them with detailed case data are something we are embarking on at the time of writing.

## References

- Alexander, C. (1979). *The timeless way of building*. New York: Oxford University Press.
- Alexander, C., Ishikawa, S., Silverstein, M., Jacobson, M., Fiksdahl-King, I., & Angel, S. (1977). *A pattern language*. New York: Oxford University Press.
- Avizienis, A., Laprie, J. C., & Randell, B. (2001). *Fundamental concepts of dependability*. Computing Science: University of Newcastle upon Tyne.
- Bateman, J. (1995). Preventive maintenance: Stand alone manufacturing compared with cellular manufacturing. *Industrial Management*, 37(1), 19.
- Bendifallah, S. A., & Scacchi, W. (1987). Understanding software maintenance work. *IEEE Transactions on Software Engineering*, SE-13(3), 311–323.
- Beyer, H., & Holzblatt, K. (1998). *Contextual design: Defining customer-centred systems*. San Francisco: Morgan Kaufmann Publishers. ISBN 1558604111.
- Bobrow, D. G., & Whalen, J. (2002). Community knowledge sharing in practice. *Reflections*, 4(2), 47–59.
- Boehm-Davis, D. A., Fox, J. E., & Philips, B. H. (1996). Techniques for exploring program comprehension, empirical studies of programmers (pp. 3–37). Washington, DC: Ablex.

- Borchers, J. (2001). A pattern approach to interaction design. *AI & Society*, 15, 359–376 (*Human Computer Interaction*, 21(1), January 2001).
- Carstensen, P. (1999). Here is the knowledge-where should I put it? Findings from a study of how knowledge spaces are used within a support group. In *IEEE 8th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET ICE '99) Proceedings*.
- Charmaz, K. (2006). *Constructing grounded theory: A practical guide through qualitative analysis*. London: Sage.
- Cooke, F. L. (2000). Implementing TPM in plant maintenance: Some organisational barriers. *International Journal of Quality & Reliability Management*, 17(9), 1003–1016.
- Crabtree, A., Rouncefield, M., & Tolmie, P. (2012). *Doing design ethnography*. London: Springer.
- Crosby, M. E., Scholtz, J., & Widenbeck, S. (2002). The roles beacons play in comprehension for novice and expert programmers. In *14th Workshop of the Psychology of Programming Interest Group, Brunel University* (pp. 58–78).
- Dearden, A. M., & Finlay, J. (2006). Pattern languages in HCI: a critical review. *Human Computer Interaction*, 21(1), 49–102.
- Denef, S. (2012). A pattern language of firefighting frontline practice to inform the design of ubiquitous computing. *Communications in Computer and Information Science*, 277, 308–312.
- Dovey, K. (1990). The pattern language and its enemies. *Design Studies*, 11, 3–9.
- Erickson, T. (2000). Lingua Francas for design: Sacred places and pattern languages DIS '00. In *Proceedings of the 3rd conference on designing interactive systems: Processes, practices, methods, and techniques* (pp. 357–368). New York: ACM.
- Gabriel, R. (1996). *Patterns of software*. New York: Oxford University Press.
- Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1993). *Design patterns: Abstraction and reuse of object-oriented design* (Lecture notes in computer science, Vol. 707, pp. 406–431). Berlin: Springer.
- Gits, C. W. (1992). Design of maintenance concepts. *International Journal of Production Economics*, 24, 217–226.
- Glaser, B. G., & Strauss, A. L. (1967). *The discovery of grounded theory: Strategies for qualitative research*. Chicago: Aldine.
- Goodyear, P. (2005). Patterns, educational design and networked learning: Patterns, pattern languages and design practice. *Australasian Journal of Educational Technology*, 21(1), 82–101.
- Hughes, J. A., Randall, D., & Shapiro, D. (1992). From ethnographic record to system design. In *Computer supported cooperative work* (Vol. 1, pp. 123–141).
- Hughes, J. A., Randall, D., & Shapiro, D. (1993). Designing with ethnography: Making work visible. *Interacting with Computers*, 5(2), 239–253.
- Hughes, J., King, V., Rodden, T., & Anderson, H. (1994). Moving out of the control room: Ethnography in system design. In *Proceedings of the ACM conference on computer supported cooperative work, Chapel Hill, North Carolina* (pp. 429–438). New York: ACM Press.
- Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.
- Kaptelinin, V., & Nardi, B. (2006). *Acting with technology: Activity theory and interaction design*. Cambridge: MIT Press.
- Kemerer, C., & Slaughter, S. (1997). Determinants of software maintenance profiles: An empirical investigation. *Software Maintenance: Research and Practice*, 9, 235–251.
- Ko, A. J., Aung, H. H., & Myers, B. A. (2005). Eliciting design requirements for maintenance-oriented IDE's: A detailed study of corrective and perfective maintenance tasks. Human-Computer Interaction Institute, Paper 179.
- Laporte, T. (1996). High reliability organizations: Unlikely, demanding and at risk. *Journal of Contingencies and Crisis Management*, 4, 60–71.
- Laporte, T. R., & Consolini, P. M. (1991). Working in practice but not in theory. *Journal of Public Administration Research and Theory*, 1(1), 19–47.

- Laporte, T., & Consolini, P. (1998a). Theoretical and operational challenges of “high-reliability organizations”: Air-traffic control and aircraft carriers. *International Journal of Public Administration*, 21(6-8), 847–852.
- Laporte, T., & Consolini, P. (1998b). Working in practice but not in theory: Theoretical challenges of “high-reliability organizations”. *Journal of Public Administration Research and Theory*, 1(1), 19–48.
- Lin, J., & Landay, J. A. (2002). Damask: A tool for early-stage design and prototyping of multi-device user interfaces. In *Proceedings of the 8th international conference on distributed multimedia systems (2002 international workshop on visual computing)*, San Francisco, CA, September 26–28, 2002 (pp. 573–580).
- Marquez, A. C. (2007). *The maintenance management framework: models and methods for complex systems maintenance*. Berlin: Springer.
- Martin, D., & Sommerville, I. (2004). Patterns of cooperative interaction: Linking ethnology and design. *ACM Transactions on Computer-Human Interaction*, 11(1), March (the patterns archive is available at <http://polo.lancs.ac.uk/patterns/>)
- McDonald, N., Corrigan, S., Daly, C., & Cromie, S. (2000). Safety management systems and safety culture in aircraft maintenance organisations. *Safety Science*, 34, 151–176.
- Nakagawa, T. (2005). *Maintenance theory of reliability*. Berlin: Springer.
- Orr, J. (1996). *Talking about machines: An ethnography of a modern job*. Ithaca: Cornell University Press.
- Orton, J. D., & Weick, K. (1990). Loosely coupled systems: A reconceptualization. *Academy of Management Review*, 15, 203.
- Perrow, C. (1999). *Normal accidents: Living with high-risk technologies*. Princeton: Princeton University Press.
- Pipek, V., & Wulf, V. (2003). Pruning the answer garden: Knowledge sharing in maintenance engineering. In *Proceedings of the 6<sup>th</sup> European Conference on Computer Supported Cooperative Work*, 14–18 September, Helsinki, Finland.
- Randall, D., Harper, R., & Rouncefield, M. (2007). *Fieldwork for design*. London: Springer.
- Randall, D., Hughes, J., & Shapiro, D. (1993). Systems development – the fourth dimension: Perspectives on the social organization of work. In P. Quintas (Ed.), *Social dimensions of systems engineering: People, processes, policies and software development* (pp. 197–214). Hemel Hempstead: Ellis Horwood.
- Raouf, A., & Ben-Daya, M. (1995). Total maintenance management: A systematic approach. *Journal of Quality in Maintenance Engineering*, 1(1), 6–14.
- Rising, L. (1998). *The patterns handbook: Techniques, strategies, and applications*. Cambridge: Cambridge University Press.
- Salingaros, N. (2000). The structure of pattern languages. *Architectural Research Quarterly*, 4, 149–162.
- Schmidt, K. (2012). The trouble with ‘tacit knowledge’. *International Journal of Computer Supported Cooperative Work*, 21(2–3), 163–225.
- Smith, R., & Hawkins, B. (2004). *Lean maintenance*. Oxford: Elsevier.
- Smith, R., & Mobley, R. K. (2008). *Rules of thumb for maintenance and reliability engineers*. Oxford: Elsevier.
- Strauss, A., & Corbin, J. (1997). *Grounded theory in practice*. London: Sage.
- Swanson, L. (2001). Linking maintenance strategies to performance. *International Journal of Production Economics*, 70, 237–244.
- Swanson, E. B. (1976). The dimensions of maintenance. In *ICSE ’76 Proceedings of the 2nd international conference on software engineering* (pp. 492–497). New York: ACM Press.
- Tan, W.-G., & Gable, G. (1998). Attitudes of maintenance personnel towards maintenance work: A comparative analysis. *Journal of Software Maintenance: Research and Practice*, 10(1), 59–74.
- van Welie, M., & van der Veer, G. (2003). *Pattern languages in interaction design: Structure and Organization Human-Computer Interaction – INTERACT’03M* (Rauterberg, M., Menozzi, M., & Wesson, J., Eds., pp. 527–534). Amsterdam: IOS Press.

- Weick, K. (1976). Educational organizations as loosely coupled systems. *Administrative Science Quarterly*, 21(1), 1–19.
- Wireman, T. (2004). *Benchmarking best practices in maintenance management*. New York: Industrial Press.
- Yamauchi, Y., Whalen, J., & Bobrow, D. (2003). Information use of service technicians in difficult cases. In 2003 Proceeding CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems.