3

Work Studies and Design

Ethnomethodological studies of work are often simply referred to as ethnographic studies in a design context. In a social science context the term ethnography delineates little more than a distinction between quantitative and qualitative methods of social research. As Shapiro (1994), commenting on the limits of ethnography in CSCW, describes matters,

> While 'ethnography' as a term strikes a useful contrast to traditional methods of requirements capture, within sociology and anthropology themselves it denotes rather little. It marks a distinction between quantitative and qualitative approaches to social science and carries with it a commitment to a period and degree of immersion in the social setting being studied that is sufficient to reach a qualitative understanding of what happens there. These are important matters, but beyond this, ethnography can be put to the service of virtually any theoretical school: there are, for example, functionalist, structuralist, interactionist, Weberian and Marxist ethnographies.

This is not the place to explore the differences between such schools of thought. It is the place, however, to note that the term ethnography denotes neither a unified method nor a coherent school of thought. Rather, and as Shapiro makes clear, the term ethnography is a gloss on various and different *analytic formats*.[19] In turning to ethnography,

[19]One such format of some prominence in HCI and CSCW alike is provided by Activity Theory (Kuutti 1996). Activity theorists often employ ethnography as a means of studying their unit of analysis – i.e., activity – in the first instance. In the second instance they seek to analyse the ethnographic data they gather in terms of the constructs from which the theory is composed: constructs such as 'actions', 'operations', 'motives', 'goals', and 'tasks', etc. In other words, activity theorists use the generic analytic format furnished by Activity Theory to code and classify the ethnographic data and thereby grab onto little bits of the real world in order to render the theory real worldly. The accounts of work practice produced through the use of Activity Theory – and any other generic analytic format employed to analyse ethnographic data – are thus subject to the problem of constructive analysis outlined in Chapter 2.

systems designers are not turning to some distinct coherent entity, then, but to a varied and competing array of analytic formats, not all of them theoretical in character. One such framework is the ethnomethodological one which has, to use a phrase of Shapiro's, 'dominated CSCW' in light of Lucy Suchman's pioneering work in the field of human–computer communication (Suchman 1987). As Shapiro puts it in accounting for the format's purchase in design,

> ethnomethodology sets for itself a strict agenda which separates it in certain ways from most mainstream social science. It insists on a rigorously descriptive rather than theoretical program, or an explanatory one (in the sense that most social sciences would understand it). This lends it its strength in producing rich descriptions of work-in-context

Ethnomethodology rejects theory in order that work-in-context may be appreciated in concrete detail and that designers may, therefore grasp "what is really going on" in the course of a piece of work, "what is really the problem" about doing it, and what instruments might therefore be devised to help resolve these problems (Hughes et al. 1992). Although the contextual and non-theoretical character of ethnography has proved to be of great value to designers (Kensing and Simonsen 1997), the approach is not without its own practical problems. As Bardram (1996) puts it,

> From the very beginning, workplace studies have played a prominent role in the research field of Computer Supported Cooperative Work. They are used to understand and shed light on work and interaction happening in a workplace ... and as such [have provided] an important insight into the subtleties of ... socially constructed work practices. Within CSCW the value of these insights into the social nature of work activities, gained through such workplace studies, is unquestionable. However, there has been an ongoing dispute in the field ... [as] to the exact value of these often very detailed and specific investigations of the workplace. Questions like: how effective is the field study approach for informing [the] design of CSCW systems? How can typical ethnographic field studies, which take months or years, be done within the fast pace of systems development? What should they be used for within the design process? Are they economical, or even practically desirable in a complex design process? Is it possible to generalize such detailed and narrow studies into applicable design recommendations? ... Questions like these are often coming from the social scientists themselves.

3.1 The Role of Ethnomethodological Studies of Work in Design

In the first instance, one of the primary roles of ethnography to date has been to *sensitize* designers to the sociality of work (e.g. Heath and Luff 1992; Hughes et al. 1994; Bowers et al. 1996) – a job which, as Bardram points out,

has been admirably achieved. The purpose of sensitizing designers to the sociality of work is not necessarily one of addressing design issues directly but rather of identifying broader issues upon which effective design turns – issues to do with what to automate and what to leave to human skill and judgement, for example. Deciding what to automate and what to leave to human beings is an important matter, not least because of the critical consequences of inappropriate design decisions. Predicated on first-hand observation, ethnographic studies of work are a primary resource in the effort to identify appropriate target areas for design. That said, and as Bardram points out, ethnographers have recognized the need for their craft to become far more responsive to the needs of designers and to the practicalities of design in particular. In its home territories ethnography is a lengthy and time consuming exercise, whereas design is a fast paced and essentially 'satisficing' activity (Shapiro 1994), doing the best it can within constraints of time, budget, and with the resources to hand. If ethnography is to be incorporated into the already diverse collection of methods, tools, and techniques used in system design, then it has to be tailored to suit the satisficing character of design. Recognition of this need has seen the development of a number of discrete yet complimentary strategies, tailoring ethnography to meet the demands of the design process. These include quick and dirty ethnography, concurrent ethnography, evaluative ethnography, and re-examination of studies (Hughes et al. 1994).

3.1.1 Some Practical Strategies for the Use of Ethnography

Quick and Dirty Ethnography

Quick and dirty studies are designed to support scoping activities in design, particularly in large-scale settings where it is important to get an overview of the work of the site in as short a time as possible (Figure 3.1). Quick and dirty studies are designed to help designers develop a 'picture' of the workplace in a relatively short time frame and, in so doing, support their strategic decision making in selecting aspects of work that are important to design. Quick and dirty studies do not set out to furnish an exhaustive view on the work of the site but instead seek to map and make visible through the assembly of instances the working division of labour, the work activities it is composed of, and the major interdependencies between activities of work. Quick and dirty studies provide a broad understanding of the work domain in a relatively short period of time (the relativity of the matter depending on the size of the Organization being studied and designed for). They provide an informed sense of what the work is like and what it consists of and may be built upon in an iterative design life cycle to provide a more focused investigation of the work of the site by undertaking concurrent ethnographic studies.

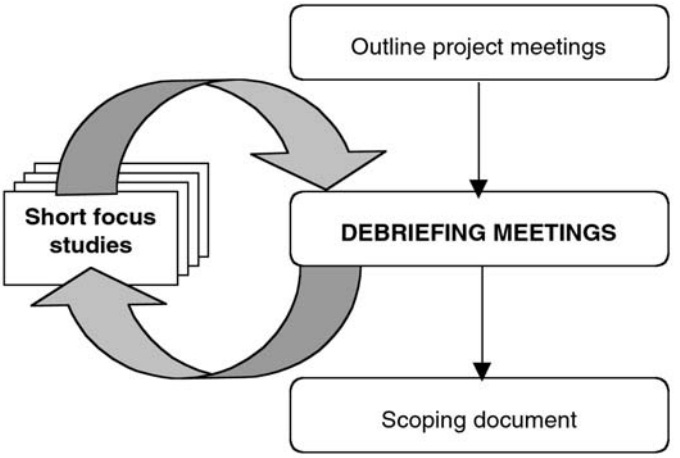

Figure 3.1 Quick and dirty ethnography (Hughes et al. 1994). © ACM, Inc. Reprinted by permission.

Concurrent Ethnography

Concurrent ethnography is a parallel process in which investigation of work and systems design proceed at the same time (Figure 3.2). Concurrent ethnography was originally configured to precede design. There is no good reason to exclude design from pursuing its own inquiries as to systems use in the first instance and the process may be adapted accordingly.

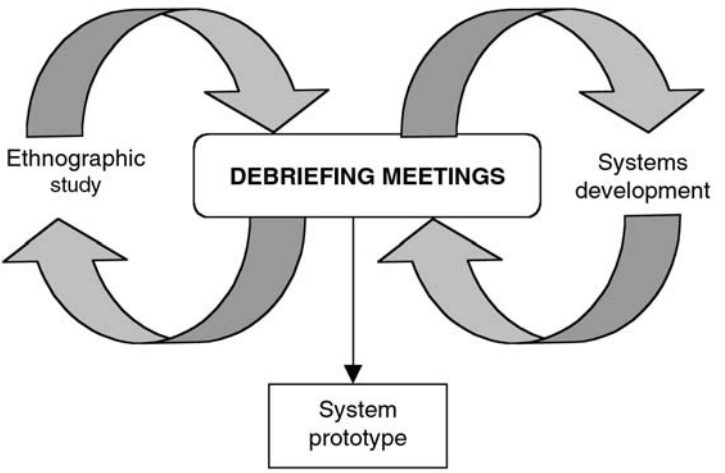

Figure 3.2 Concurrent ethnography (Hughes et al. 1994). © ACM, Inc. Reprinted by permission.

Following initial investigations, a debriefing meeting is held between the ethnographer and the designers in order to identify, discuss, and elaborate issues of relevance to design. With the findings in hand, iteration of the prototype may proceed and further fieldwork may be undertaken with respect to the designers' concerns. Concurrent ethnography is a directed process, then, in which each stage of the fieldwork is intended to target issues raised by designers in debriefings. It is a process characterized by fieldwork→debreifing→prototype-iteration→fieldwork. Concurrent ethnography is a very flexible process which may be undertaken as and when required and for as many iterations as required. Having developed a working prototype, evaluative ethnography may then be undertaken.

Evaluative Ethnography

Evaluative ethnography is a more focused version of quick and dirty ethnography, which is to say that it does not require a prolonged period of fieldwork (Figure 3.3). The purpose of evaluative ethnography is provide a 'sanity check' of design proposals or of an existing prototype, where analytic emphasis is placed on establishing the 'work-ability' of the proposed design solution – i.e., to assess the proposed solution's efficacy in relation to the actual performance of work. Evaluation studies are not intended to be exhaustive but rather, to assess the prima facie viability of technical solutions as seen and understood from the point of view of

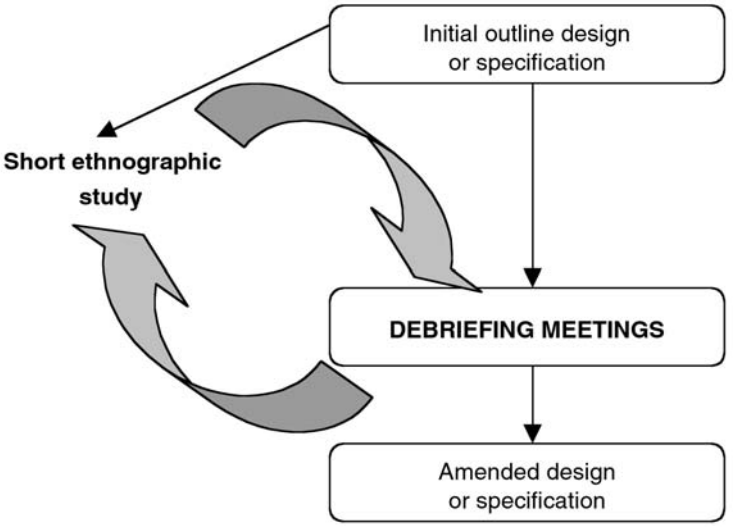

Figure 3.3 Evaluative ethnography (Hughes et al. 1994). © ACM, Inc. Reprinted by permission.

work accomplishment and to identify problematic issues that may subsequently emerge. Evaluation studies may also be extended to support more rigorous assessments of a developed prototype (an issue that will be addressed in due course).

Re-examination of Studies

This is an entirely different order of ethnographic usage, emergent from wider research endeavours. The purpose here is to assemble a corpus of studies that may be drawn upon to identify common issues across a variety of domains. Such findings may be used to sensitize designers to common arrangements of cooperation in various workplaces and to issues that may therefore be of relevance in particular kinds of design undertaking. Experience in the field also suggests that it may be useful to re-examine studies as and when particular problems in a design project emerge. Being a satisficing activity, oversights are made, aspects of work are disregarded, certain issues are put on hold or ruled out of scope for the time being, and so on, and re-examining previous studies conducted in the course of a project can, on occasion, be a productive way of addressing obstinate design problems.

These four strategies should not be understood as laying out mutually exclusive or rigidly demarcated approaches to the study and analysis of cooperative work in design. On the contrary, the four modes of investigation are intended to complement and shade into one another as design work demands. In contrast to ethnography's employment in a social science context, none of the above approaches require prolonged periods of immersion in the field. While this may be disconcerting to social science purists, it should be understood and taken seriously that the point of work study in design is not to conduct social science research for its own sake. Rather, the point is to develop social science 'methods' (i.e. analytic perspectives) in order that designers may develop an awareness of the real world character of work and do so in a manner that fits in with and enhances *their* day-to-day working practices. As Hughes et al. (1994) remind us here,

> much can be learned from relatively short periods of fieldwork. Indeed, within the context of design, and we emphasise this, diminishing returns to fieldwork set in relatively quickly. In other words, fieldwork of prolonged duration is not always necessary in that it would be more effective to direct that effort in accordance with design objectives.

The turn to the social, and ethnomethodologically-informed ethnography in particular, is not in itself a turn to prolonged courses of social research in design, but to adapted social science 'methods' to provide perspicuous representations of the sociality of work within the constraints

of, and for the for purposes of, design. Ethnography, as Sharrock (2000) points out, has no strategic role to play in design. It offers no radical transformations of design practice (as some have promised), but acts as the designer's eyes and ears on the ground. The practical strategies outlined above represent primary configurations of an empirical approach to describing, analysing and representing the sociality of work *within the flow of design work*. Although these configurations have proved to be effective ways of weaving ethnographic insights into the design process, they have not passed without complaint from some quarters, however.

Bardram (1996) suggests, for example, that the ethnographic strategies outlined above create a problem of

> one-way communication between users and designers, meaning that information is floating from the work practices to the designer, but no information about the future technology, the use of computers etc., is floating back to the future users in the workplace.

This is not entirely true as ethnographers often act as communicative agents between users and designers in the course of conducting fieldwork. Telling users why they are being studied, for what reasons, and towards what ends, is an essential part of getting access to the worksite and getting permission to study the work. Neither is the problem of one-way communication a problem for design per se but for Participatory Design (Schuler and Namioka 1993). Participatory Design is dedicated to involving users directly and actively in the process of design. While ethnography has no objection to direct user involvement in the design process, the economic realities of design often mitigate against this, as leading advocates of the Participatory Design movement are well aware (Grønbæk et al. 1997).[20] Nonetheless, the problem of one-way communication, and thus the separation of users and designers, is a significant problem to be reckoned with and it is in this respect that ethnography has been charged (quite rightly) with misrepresenting itself as a 'proxy user' (Kyng 1995). Such criticisms come in light of comments to the effect that ethnographers may act as 'users champions' in the early stages of design (Bentley et al. 1992). More considered reflections on the ethnographer's role in the design process recognize that the ethnographer can never know the work domain as users know it (Hughes et al. 1993). However, it is not (or should not be) the ethnographer's responsibility to speak on behalf of or represent users in the political sense of the word (and Participatory Designers are deeply concerned with the politics of representation, as will be addressed in due course). The ethnographer's

[20]Information about the use practice is not, after all, a free product in design and user involvement can be a costly exercise (especially when contrasted with the relatively inexpensive costs of ethnographic inquiry).

responsibility is, instead and quite legitimately, one of representing the sociality of the job in the descriptive and ethical sense of the word (ethical in the sense of being faithful to phenomena and thus describing it in a practically adequate fashion, and in the sense of not putting the persons observed at risk of sanction, etc.). Representing the sociality of work requires the ethnographer to explicate the work practices whereby the coordination of workaday activities gets done, day in and day out and by different production cohorts or members of staff. Thus, the ethnographer is concerned with portraying those features of work that tie individuals together, which maintain regardless of particular individuals, and which individuals may, therefore, find it very difficult to articulate (Crabtree 1998). Knowledge of the cooperative work practices through which work is organized by parties to the work, whether those parties be co-located or distributed across space and time, is what the ethnographer brings to design. The ethnographic methods and strategies outlined above provide tried and tested methods for accomplishing this goal.

The effort to bring an awareness of the cooperative aspects of work to bear on design has prompted further complaints from design practitioners regarding the *link* between studies of work and design. As one computer scientist describes matters here,

> Look, I don't care about all this ethno babble, just tell me what to build! All you ever do is tell me stories. Why is this useful? What can I do with it? (Cited in Crabtree et al. 1997)

Ethnographers in the field originally responded to the challenge in one of two ways, attempting to address the requirements problem on the one hand and contesting ownership of the problem on the other, arguing that it was and is no part of ethnography's remit to come up with design solutions. Ethnography's role is to 'impart knowledge' as to the cooperative work of intended users, not to 'give form' to potential design solutions supporting that work (Plowman et al. 1995).

It might be said that those who contested ownership of the problem were right to do so only in the sense that ethnographers rarely possess the necessary technical competence to formulate design solutions on their own. After all, what constitutes a solution very much depends on the technology being developed, on the purpose of design, on the often divergent needs of users, and on a host of other contingent Organizational factors. The ethnographer cannot sit down, then, and analyse a work study with an eye towards formulating design solutions, as just what constitutes a solution will be the outcome of considering a host of contingent factors (social, technical, economic, political, and the rest). This, however, does not rule out a constructive role for ethnography in the formulation of design solutions. Adopting any such role will require that ethnographers move beyond 'imparting knowledge' to directly inform the construction

of design solutions in collaboration with the other parties to systems development. Indeed, it could be argued with good reason that satisfying that requirement is essential to the approach's long-term inclusion in design. As Shapiro (1993) puts it,

> Any role at all for sociologists [anthropologists or ethnographic work analysts in general] in this field rests on their claim to being in a better position to identify particular aspects of 'what is really going on' in a given field of work and 'what is really the problem' that people encounter in doing it. If this claim is not sustainable then sociologists [etc.] have no contribution to make to systems design.

If ethnography cannot actively and constructively support system developers in the reorganization of work through technology design rather than run for cover, then it has no business in the field. While rightly leaving the actual production of design solutions to designers, ethnographers are nevertheless obliged to assume a constructive role in design if their craft is to be of any lasting utility. This means that the ethnographer will have to rid themself of disciplinary baggage and become a *bricoleur* in design practice.

3.2 Using Ethnography to Give Form to Design (The Bricoleur's Craft)

The problem of linking ethnography to design and thus of giving form to design is primarily a problem of figuring out what work studies mean in a design context. As ethnographers rarely possess technical competence, working out what work studies mean for design requires direct collaboration with designers, involving them in cooperative analysis of studies of the design space and identification of user needs. Thus, the problem becomes one of communication – that is, of conveying the findings of work studies to designers and in readily digestible ways that support cooperative analysis of the design space. Early work in the field saw the development of the Designer's Note Pad (DNP) as a means of promoting effective communication between ethnographers and designers (Hughes et al. 2000). DNP is a flexible hypertext system that supports the rapid construction of directed graphs used in structured design methods. The tool allows ethnographers to present and communicate the results of work studies in a way that fits current design practice. The presentation of findings is organized through a number of 'viewpoints' (Kontonya and Sommerville 1992) – specifically, the ecology of work viewpoint, the workflow viewpoint, and views of work viewpoint (Hughes et al. 2000).

The Ecology of Work Viewpoint

This viewpoint is concerned with representing the physical layout of the workplace and the spatial distribution of parties to the work. The purpose of this viewpoint is to furnish a sense of the physical setting within which the work takes place, the spatially situated arrangements of cooperation and assistance that obtain between parties to the work, and the artefacts employed in getting the work done.

The Workflow Viewpoint

This viewpoint is concerned with the sequential relationships that hold between workaday activities, with what the workaday activities are in a setting and how they connect together to form distinct interrelated processes through the production and transformation of information over the course of the assembly and coordination of work. The workflow viewpoint represents distinct process of work and embedded information processes.

The Views of Work Viewpoint

This viewpoint is a collection of viewpoints on work, specifically, distributed coordination, awareness of work, and plans and procedures.

- *Distributed coordination* is concerned with the ways in which the component activities of the workflow are coordinated by parties to the work. It is concerned with how *this* group of people put *this* particular job of work together and coordinate it with the next job in the flow of work.
- *Awareness of work* is concerned with the ways in which parties to the assembly and coordination of workaday activities make one another aware of the work they are engaged in and its status. It is concerned with the ways in which 'where in the process we are now' and 'what needs to be done next' is displayed by and for parties to the work in order that they might coordinate their efforts effectively.
- *Plans and procedures* is concerned with the specification of Organizational requirements and controls. In paying attention to distributed coordination and awareness of work, emphasis is placed on the ways in which – i.e., on the cooperative work practices through the accomplishment of which – Organizational requirements are actually met.

Thus, the views of work viewpoint is concerned with ways in which the workflow is actually produced in the real world, real time actions and interactions of parties to the work. Each viewpoint is presented in diagrammatic form and linked to relevant viewpoints. Thus, and for exam-

ple, the plans and procedures governing a particular process of work may be linked to the workflow diagram, which may be linked to the distributed coordination and awareness diagrams, and the ecology diagram. In the linking of viewpoints, ethnographers may articulate the real world, real time character of work and embedded information processes, the activities involved, the cooperation they entail, and the ways in which artefacts are used to get the work done. The work of articulation is supported through textual annotations attached to the viewpoint diagrams and/or objects therein. These may describe the real world work of particular workers assembling and coordinating particular jobs and/or the use of particular artefacts (computers, documents, files, etc.) embedded within the workplace, as observed by the ethnographer.

Attempts to marry viewpoints with industry standard notation, specifically the Unified Modelling Language or UML (Fowler and Scott 1997) have seen the development of the Coherence Method (Viller and Sommerville 1999). The intention here is to enable designers to extract systems requirements from the findings of work studies. The method is predicated on Jacobson's (1995) use-case driven approach. Use-cases define the work a system should support by describing the different users or more specifically 'actor nodes' of a system, whose work is defined through a number of 'use-case nodes' which combine to form a use-case model. Actors exist outside the system, whereas use-cases exist inside the system. Actor nodes are connected to use-case nodes through 'communication arcs', which allow stimuli to be sent between 'instances' of the actor classes and use-cases classes that the model is composed of. Instances of actor classes are *sequences of actions* engaged in by users in the course of performing particular activities, for example, withdrawing money from a cash dispenser. Instances of use-case classes are sequences of actions produced by the machine in response to user actions, for example, instructing users to input a PIN code, to specify the amount of money to be withdrawn, and checking the user's balance before issuing money,. Thus, use-cases are sequenced design solutions to sequences of user actions and interactions. Use-cases are envisioned courses of machine response to concrete user actions performed on and through a system. They are produced through observations of work, which are used to sketch out potential design solutions. As Jacobson describes matters,

> Use-case modelling is an important tool to develop an outside view of a system. Combined with other techniques for human–computer interaction, it will help us build systems with high usability. Furthermore, it is a powerful means of communication between different kinds of people participating in the development work, such as orderers, users, requirements analysts, designers, and testers.

Formulated use-case models are subsequently employed to construct a formal analysis model, a design model, an implementation model, and a

testing model (Jacobson et al. 1992). The Coherence Method takes the notion of use-case modelling, combines it with the viewpoints approach and applies it to the ethnographic record. Modelling proceeds from the plans and procedures viewpoint (i.e., from the Organizational viewpoint) but describes observable-reportable workflows, collaborations, and the artefacts used in terms of sequence diagrams, thus reconciling formal Organizational requirements with the actual circumstances of work. The production of sequence diagrams enables designers to identify sequential orders of action that are the focus of support in design (see Figure 3.4, for example). Having identified sequential orders of work, designers may then set about formulating use-cases to support those sequences of practical action.

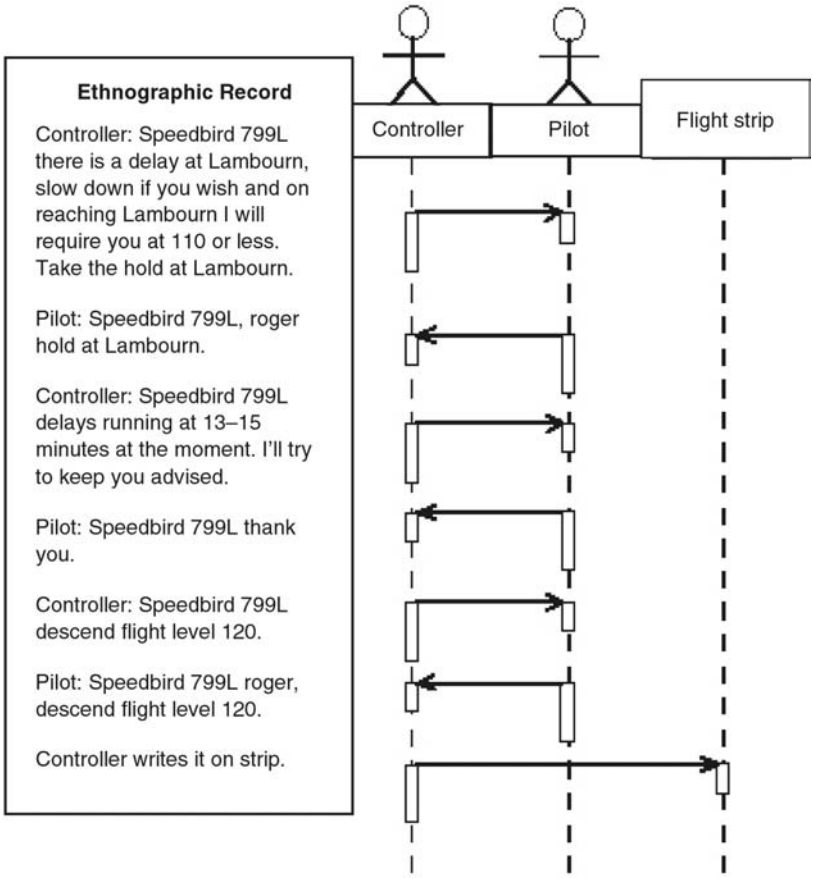

Figure 3.4 Identifying sequences of action (Viller and Sommerville 1999). © 1999, IEEE. Reprinted by permission.

Despite its appeal to the engineering mentality, the Coherence Method is not without its problems. As Viller and Sommerville (1999) put it,

> Whilst it could be claimed that Coherence has removed the communication problem between sociologists and software engineers, the effect has been to shift the problem from one of understanding people in conversation, to understanding people via the method's documentation.

While there is much of value to viewpoints and the use-case driven approach, the Coherence Method does not so much solve the problem of communication as side-step it. The ethnographer's job finishes with the structuring of the ethnographic record in terms of viewpoints, and requirements are identified by designers through the application of standard notation to that organized ensemble of data. This assumes that designers will see the data in the same way as an ethnographer – that designers will appreciate the data's significance in a similar fashion. Unfortunately this not the case and there is, then, a serious risk that although informed by the real world, real time organization of cooperative work, designers will formulate perfect technical solutions to the *wrong problems* of work.

Take the design situation of computer support for Air Traffic Control (ATC) represented in Figure 3.4, for example. While design would be right to focus on the cooperative work of the controller and the pilot, attending in particular to the sequences of action that take place between the two, it would be a mistake to think that those sequences identify targets and requirements for computer support. One might, for example, target the resources used (e.g. the flight strips) and take requirements to be articulated by the sequences of actions implicated in the use of those resources (e.g. in the cooperative work of the controller and the pilot). If we consider the use of the flight strip, a piece of paper about one inch wide and eight inches long, a rather different picture emerges from an ethnographic point of view, however.

The flight strip is an essential resource in the work of controllers, work that might be described as "manoeuvring aircraft in an ordered way through the skies" (Hughes et al. 1988). Seen from a controller's point of view, manoeuvring aircraft in an ordered way through the skies first requires that a picture of the skies is available to work with. 'Getting the picture' is of paramount concern to air traffic controllers and it might be thought that this is a relatively straightforward matter of consulting the radar. In practice, however, radar is not sufficient for getting a picture of the skies.

> (Fieldnote extract) It would be an impossible job to sit and look at the radar and look at all the different blips and try to avoid [conflictions] by putting the aircraft into blank spaces on the radar, so you've got to have [some] other information to tell you what traffic is coming into and going out of your sector. (Air Traffic Controller)

That other information is furnished by the flight strip, which provides information concerning the flight path of each individual aircraft.

> (Fieldnote extract) From your strips you can find out whether or not there is a possible confliction and what you can do about it. You then go to your radar and look for that particular aircraft and see where it is in reference to the outbound from Heathrow, for example ... Then you decide what you're going to do with it, whether you are going to go up underneath it, whether you are going to wait until its gone past, whether, if its on your frequency, whether you can put them on parallel headings and then you can climb it up to the other aircraft levels. Same as with the inbounds, it's the same sort of thing.

Information on the flight strip includes the aircraft name or call sign (e.g. Speedbird 799L), its departure and destination point, its preferred route and height, the type of aircraft and its speed. The aircraft's estimated time of arrival at particular navigation points is printed on the left-hand side of the strip beside the abbreviation for each navigation point. Each flight sector has three or four key navigation points and strips are printed and placed in racks or bays next to the controller's radar screen some ten minutes or so before each aircraft's arrival into each control sector. The flight strip represents an aircraft at each stage of its journey through a flight sector and, as such, enables the controller to get a picture of the skies *and* to work that picture, thereby manoeuvring aircraft in an orderly way through the skies.

As Hughes et al. (1992) point out, the above is not to say that the strips determine the sequence of actions undertaken by controllers in working the picture in the sense that what comes along a production line determines what a line worker has to do next.

> Rather, the controller has to organize the strips so that they can become an instrument that helps organize, and so make possible, controlling work. Strips are 'manipulated', 'glanced at', 'taken heed of', 'ignored', 'revised', and so on. And not just when they are first placed on the racks, but continuously all the time they are in use. The end result of these activities is that, at any moment in time, what the strips indicate and create [in their use] is the sequence of controller actions that results in the achievement of order in the skies. Thus, management of the strips constitutes a large part of the work that underscores controlling competence.

Managing or working the strips is the cooperative achievement not only of the controller and the pilot but also, *of the ATC team*. When a strip is first printed (information is initially generated from flight plans submitted by the pilot and augmented with weather details), the controller's assistant places it in the appropriate rack. Placing the flight strip is not a mechanical task but requires that the assistant check the strip for its accuracy. Errors are not infrequent – the designated flight path may be inconsistent with the flight number, for example – and it is the assistant's

responsibility to rectify any visible errors on the strip, which may well require collaboration with others outside the control room as contingencies dictate. Having tracked and changed any errors, the assistant may place the annotated strip or a newly generated strip in the flight rack.

Insofar as the strip is annotated by the assistant then annotations are indicated by the use of a distinctly coloured pen. Each member of the control team may annotate the flight strip at any time and each has a different coloured pen so that the person who made the annotations is immediately identifiable and may be quickly consulted if need be. In placing the strip on the rack, the assistant does not just put it anywhere. The rack is organized in terms of two distinct sections: pending and live. Furthermore, as there is a temporal flow of aircraft into each sector, so in placing strips into the pending rack the assistant organizes them in terms of time sequence. More still, in organizing the temporal placement of strips, problems may become apparent – it may be seen that two aircraft are projected to reach the same navigation point simultaneously and at the same height, for example. On recognizing a potential confliction through the action of placing the strips, the assistant or any other member of the team indicates the problem by cocking the problematic strips out of the rack. This makes it immediately obvious to the controller that when the time comes to deal with those flights a problem will have to be considered and, as Hughes et al. (1992) note, to the practised eye it will be obvious at a glance what the problem is.

Cooperative work on the strips continues once they become live. A strip usually becomes live on receipt of a radio message from the aircraft crew as it enters the sector. On receiving this message, the controller removes the strip from the pending rack and places it in the live section. Once again, placement of the strip is a highly organized affair. In placing the strip in the live section, the controller places it in that part of the rack designated by the navigation point being used. Furthermore, just as the assistant orders strips according to time sequence, then so too does the controller. Thus, as strips become live a temporal flow emerges, with new strips at the top and old strips at the bottom (or vice versa as the controller prefers) and from navigation point to navigation point. More still, as and when strips become live they are colour marked by the controller to indicate that they have moved them. In the course of being live, strips may be replaced by the sector chief or the assistant with more up-to-date or revised ones. Such replacements lack the controller's mark and the controller can, then, see at a glance that the strip requires attention and that its implications be taken on board. The controller also marks any commands issued to the aircraft on the live strip (as recognized in Figure 3.4). This provides a publicly available record of just what the controller has done to date. Finally, as and when the aircraft crosses each navigation point, the controller crosses through it on the flight strip. This indicates the last point through which the aircraft passed in the sector. When the

aircraft crosses the last navigation point in the sector, the controller directs the aircraft to contact the next sector controller and then crosses through the strip. This indicates that the controller's work has been completed.

Seen from a design perspective, the sequences of action implicated in organizing the flight strips might identify requirements for a computer-based artefact replacing the paper strips. That artefact might display automatically generated flight strip information, annotate changes automatically, automatically order flight information in pending and live status, display controller's marks and the other actions that comprise the discrete sequences implicated in the working of the strips. However, automation of the flight strips overlooks the fact that the sequences of action that the working of the strips is composed of are themselves embedded in and produced through the cooperative work of the ATC team. This is not a trivial point. As Hughes et al. (1992) inform us, ATC in the UK has the best record in the world – not a single civil collision has ever been attributed to air traffic control.

> If one looks to see what constitutes this reliability, it cannot be found in any single element of the system. It is certainly not to be found in the equipment [radar, etc.] ... Nor is it to be found in the rules and procedures, which are a resource for safe operation but which can never cover every circumstance and condition. Nor is it to be found in the personnel who, though very highly skilled, motivated and dedicated, are as prone as people everywhere to human error. Rather, we believe it is to found in the cooperative activities of controllers across the 'totality' of the system, and in particular in the way it enforces the active engagement of controllers, chiefs and assistants with the material which they are using and with each other. This constitutes a continuing check on their own work and a cross-check on that of others.

It might otherwise be said that the reliability of work in the air traffic control suite is to be found in the cooperative work practices of its staff. Those work practices are designed by staff to construct a protective cocoon around the controller and in so doing to provide all the resources required so that the controller can concentrate exclusively on the work of ordering the skies. That cocoon is not an accidental construct nor is it a product of a simple division of labour and allocation of tasks. On the contrary, it is an active construct created *moment by moment* through the sequences of action embedded in and produced through the cooperative work of the ATC team organizing the flight strips. Ordering the skies in a safe manner is the cooperative accomplishment of the ATC team's collective working of the flight strips. Thus, when considering the social organization of sequences of action, it is important to consider one sequence's relation to another (such as the relation of pilot–controller interaction to controller–assistant interaction). As Viller and Sommerville (1999) point out, there is no guarantee that the Coherence Method will be sensitive to

the social relationships and constellations of cooperation and assistance that shape the work of the setting. This means that although designers may be able to identify much of relevance to design through the Coherence Method's documentation, the work analyst cannot be dispensed with in the 'creative process of design' (Twidale et al. 1993) – a situation which again raises the problem of communication.

3.2.1 A Lingua Franca for Design

Recent efforts in the field have been directed towards establishing a *lingua franca* supporting communication between ethnographers, designers, and others involved in the development process.

> **lingua franca. 1** a language adopted as a common language between speakers whose native languages are different. **2** a system for mutual understanding. (The Concise Oxford Dictionary)

Underpinning the call for a lingua franca in design is the recognition that ethnographers, designers, and other parties to design conduct their daily business in different professional languages, a situation which compounds the problem of communication (Crabtree et al. 2000; Hughes et al. 2000). Researchers have consequently emphasized the need to develop a 'pattern language' for structuring, presenting and analysing the findings of work studies (Erickson 2000a). A wide variety of pattern languages have been developed for various purposes in design, including those of software engineering (Gamma et al. 1995), HCI (Bayle et al. 1998), and work study (Erickson 2000b). Software engineering and HCI pattern languages are marked by a concern to address problems encountered in previous development projects and to make solutions available to the wider community of practitioners. Pattern languages constructed to support work study eschew the problem–solution focus and place emphasis instead on the description of common arrangements of cooperative work (Martin et al. 2001). Such languages are essentially *retrospective* in character, re-examining previous ethnographic studies in order to discover patterns of interaction or cooperation that occur across a variety of settings.

Recent developments in the field have sought to develop a prospective pattern language to support communication and cooperative analysis of the design space (Crabtree et al. 2001b; Crabtree and Rodden 2002). This approach seeks to use a pattern language as a means of conveying the important details of ongoing work studies to designers in the flow of design work. Like all other pattern languages in design, this prospective approach draws on the seminal work of the architect Christopher Alexander (1979). It is also inspired by existing design practices that are

central to the analysis of the design space. Specifically, the prospective approach is predicated on Sharrock and Anderson's (1994) study of how analysis of the design space and user needs is structured in design practice through the commonsense method of typification (as discussed in Chapter 1). The method consists of the appeal to and use of knowledge of patterns of action to provide for cooperative analysis of user needs and to construct an inter-subjective understanding of the design space.

The commonsense method of typification presupposes the use of formal methods in design and may be supported through the adaptation of Alexander's original architectural framework to structure and present ethnographic findings to designers. Alexander's framework is constructed on the foundational observation that towns and buildings are organized through reoccurring patterns of action (or "patterns of events" in his own terminology) that people take part in over and over again. Being in bed, taking a shower, making breakfast, sitting in the study writing, walking in the garden, eating lunch, going to the movies, taking the family to a restaurant, having a drink at a friend's house, driving on the highway, and going to bed again are examples used by Alexander to illuminate the point. Our lives are organized through *reoccurring patterns* of work, leisure, travel, relaxation, and the rest.

While patterns of action are implicated in the daily lives of individuals, a great many patterns of action are not individualistic but organize our lives together as members of society:

> they are the rules through which our culture maintains itself, keeps itself alive, and it is by building our lives out of these patterns of events, that we are people of our culture.

Thus, and for example, each morning people get up, shower, eat breakfast, and drive down the highway to work, where they together engage in other patterns of action, such as checking their mail, attending meetings, or going for lunch, etc. A great many of the patterns of events whereby towns and buildings are organized are thoroughly social in character, then, and so make the social organization of towns and buildings visible.

Alexander also observes that patterns of action are *tied to particular places* within a society. So taking a shower in a morning is tied to the bathroom, eating breakfast to the kitchen, driving to work to the highway (and not the sidewalk), for example. As Alexander puts it, patterns are always anchored in space:

> I cannot imagine any pattern without imagining a place where it is happening.

The patterns of action out of which any particular place – a bathroom, a kitchen, a motorway, a library help desk, etc. – is made up are also rather small or finite and defined by social convention, which provides for the

generalization of the patterns that occur in a place to other such places in the wider society. Thus, by virtue of social convention, showering may be generalized to bathrooms, eating breakfast to kitchens, driving at high speed to motorways, intermediated searching to library help desks, and so on. Alexander's framework ties patterns of action to the architectural environments in which they naturally occur, then, and provides a basis for analysing settings and improving towns and buildings

The work of analysis consists of locating the patterns of action that occur in and across the distinct places that make up a particular setting and of explicating the patterns of relationships that exist between patterns of action and the material arrangements of place: between cars and pedestrians crossing the road, between a person entering a building and the physical entrance, between cooking and the physical layout of the kitchen, etc. (Alexander et al. 1977). Placing analytic emphasis on *patterns of relationships*, Alexander draws our attention to the reoccurring ways in which people interact with their architectural environment and, particularly, with the material arrangements that it is made up of. These patterns of relationships are the primary object of pattern analysis in Alexander's framework. They elaborate the socially organized ways in which people *use* the material arrangements of place and the problems they encounter in the course of use.

Knowledge of patterns and problems is made publicly available, along with proposed architectural solutions, through the use of a distinct presentation format, which Alexander describes in the following way.

> First, there is a [pattern number, a title and a] picture, which shows an archetypal example of [a] pattern. Second, after the picture, each pattern has an introductory paragraph, which sets the context for the pattern, by explaining how it helps to complete certain larger patterns. Then there are three diamonds to mark the beginning of the problem. After the diamonds there is a headline, in bold type. This headline gives the essence of the problem in one or two sentences. After the headline comes the body of the problem. This is the longest section. It describes the empirical background of the pattern, the evidence for its validity, the range of different ways the pattern can be manifested in a building, and so on. Then, again in bold type, like the headline, is the solution – the heart of the pattern – which describes the field of physical and social relationships which are required to solve the stated problem, in the stated context. This solution is always stated in the form of an instruction – so that you know exactly what you need to do, to build the pattern. Then, after the solution, there is a diagram, which shows the solution in the form of a diagram, with labels to indicate its main components. After the diagram, another three diamonds, to show that the main body of the pattern is finished. And finally, after the diamonds there is a paragraph which ties the pattern to all those smaller patterns in the language which are need to complete this pattern, to embellish it, to fill it out. (Alexander et al. 1977)

In talking of larger patterns Alexander may be understood to be talking about the *primary* patterns that define place – e.g. making breakfast,

lunch or dinner in the kitchen – and in talking of smaller patterns he may be understood to be talking about the *component* patterns that make up a primary pattern – e.g. getting foodstuffs from the refrigerator, implements from drawers, using the cooker or microwave, cleaning the table, etc. It is in this respect that an adapted patterns framework and presentation format might be of utility to design. Illuminating primary and component patterns of events and technological (rather architectural) relationships identified in ethnographic studies, an adapted framework might support the commonsense method of typification that underpins analysis of the design space.

3.2.2 The Adapted Patterns Framework

Refocusing Alexander's framework to address technological rather than architectural arrangements of place, allows ethnographers and designers to identify generic (or typical) patterns of relationships that obtain between people and computers in particular places (such as the library help desk) from out of the minutia of ethnographic studies – patterns that may be employed to support the formulation of design solutions. Generalization might be supported by adapting Alexander's pattern format to a web-based presentation vehicle. A web-based format means that much richer resources may be provided to designers than can be contained in a text and so the 'archetypal picture' may be replaced with actual video footage that displays the pattern in question (Figure 3.5). The video may be viewed via a hyperlink embedded in the pattern's title. The title consists of a commonsense description (e.g. searching at the help desk). Given the analytic emphasis placed on the relationship of action to technology, a sub-heading *key technologies* may be added to the format, listing keywords that describe the technologies used in the video sequence (e.g. Online Public Access Catalogue, reading list, notes, etc.).

The 'introductory paragraph' is renamed *interactional setting of the pattern* (Figure 3.6). This section briefly describes a) where the pattern occurs; b) who is involved in the sequence of interaction; c) what the parties to the interaction are doing; and d) the primary pattern this pattern is a component of (e.g. doing an intermediated search, where just what is wanted is not known in advance). Hyperlinks in this section connect the pattern to a primary *patterns log*, where access to the corpus of component patterns making up the particular primary pattern in question (e.g. intermediated searching) is provided (component patterns in intermediated searching would include user–subject librarian interaction around OPAC and the collaborative use of the physical catalogue, for example).

The 'essence of the problem' is renamed the *organizational context* and provides a formal summary of the practical issue addressed by the pattern (Figure 3.7).

Pattern No. 5 Searching at the Help Desk

[click on hyperlink to view video]

Key Technologies: Online Public Access Catalogue

Figure 3.5 The 'archetypal picture'.

Interactional Setting of the Pattern

The sequence of interaction on the video takes place at the library help desk, where users often turn when they cannot find items satisfying their information requirements. Two users are involved in the interaction along with a member of staff. Together they try to establish the identity of specific items that will satisfy the users' information requirements. The pattern of action is part of the primary pattern 'intermediated searching' [hyperlink to component patterns listed in patterns log], where just what is wanted is not known in advance.

Figure 3.6 The 'introductory paragraph'.

Organizational Context of the Pattern

Library users often turn to the help desk when they cannot find information satisfying their information requirements Help desk staff characterize this work with users as 'getting details out of people', as 'trying to find what they are looking for', or more generally and formally, as 'filtering work'. Filtering work is concerned with making user inquiries intelligible in terms of the library catalogue's organization in order that useful materials satisfying users needs might be located. Users and help desk staff conduct filtering work together through performing a course of categorization work in which they work up a concrete sense of just what is being searched for.

Figure 3.7 The 'essence of the problem'.

The 'body of the problem' is renamed *the work of the pattern* and describes the routine activities that make up the pattern. The work is summed up in a synopsis, which is followed by a transcript of the talk and description of the relevant non-verbal practical actions of the parties to the interaction (see Figure 3.8).

◆◆◆

The Work of the Pattern

Synopsis: Two users approach the help desk and ask staff where a certain section of the library is. Staff directs them to OPAC and asks them if they know how to use it. Staff takes them over to OPAC and the users explain that what it is they are looking for. Their explanation is in the form preliminary information categories, which staff uses to initiate a search on OPAC. The three of them employ the OPAC retrieval list to work up more precise information categories. This enables the users and the librarian to focus down on a manageable and sufficiently small collection of information in the catalogue and from that point, to identify and extract information of personal relevance to the user. Staff takes the users to a section of the library and locates materials the users are looking for.

Transcript:

1. Sarah: Could you tell us where market – what was it – market intelligence?
2. Lisa: Yeah.
3. Sarah: Market intelligence
4. Sylvia: Marketing is C floor. (Points to OPAC located at help desk.) Do you know how to use the screens?
5. Lisa: Yeah but
6. Sylvia: You need to find the classmark for the book.

Sylvia leaves the help desk, leads the two users (Lisa and Sarah) to a free OPAC terminal nearby and initiates a title search.

7. Lisa: It's not a book.
8. Sarah: It's like information, information about these particular products and services. It's called market intelligence and leisure intelligence et cetera et cetera.
9. Sylvia: And is that the name of
10. Sarah: That's the name – market intelligence and leisure intelligence. It's not a book as such. It's usually in the reference library.
11. Sylvia: Is, is it a serial?
12. Lisa: Yeah.
13. Sylvia: It's a serial.

Sylvia initiates a serial search on OPAC.

14. Lisa: It's a journal.
15. Sarah: It's not so much a journal but it does come out every few months.

Sylvia browses the serial search retrieval list.

16. Sylvia: Is it marketing intelligence and planning? Is that the one?

Sylvia points to an item on the retrieval list.

17. Sylvia: T6 – it's a journal.
18. Sarah: No. It's not a journal.
19. Sylvia: Do you want to check at that and find the journal itself?

Sylvia points to the item's classmark on the OPAC screen.

20. Sarah: Been there.

Figure 3.8 The 'body of the problem'.

The 'solution' is renamed *the practices ordering the work of the pattern* and articulates the recognizable social practices implicated in the work's routine accomplishment (Figure 3.9). This section describes the familiar, recurring ways in which routine workaday activities get done. Description of these reoccurring practices provides for the identification of the pattern.

The 'diagram' is replaced with a more appropriate category to the task at hand, namely *the pattern of technology usage* (Figure 3.10). This section of the format describes the technologies used in the sequence of

The Practices Ordering the Work of the Pattern

1. Formulating a specifically vague description. Users initially provide a specifically vague description of their information requirements or 'keyword' (such as 'market intelligence'). Such opening descriptions are very vague insofar as they cover many things but at the same time, and without contradiction, are also very specific as the information required is, in some yet to be articulated way, connected to 'marketing'. Library users furnish help desk staff with such descriptions as a matter of course, and so broadly circumscribe the search area.

2. Formulating preliminary information requirement categories. The connection between the circumscribed search area (e.g. marketing) and the information requirement, which is (in part) in the user's head, needs to be articulated. In articulating that connection, vague descriptions are transformed into more precise descriptions that fit the catalogue. The transformation of vague descriptions is accomplished through the cooperative formulation of preliminary information requirement categories, which consists of associating vague descriptions with categories that the library catalogue is composed of. The work of association is often accomplished through the use of OPAC and allows users and staff to work up candidate categories of solution to the search problem – that the user's information requirements concerning marketing are located in a journal, for example - thereby narrowing down the search area.

3. Formulating specific information requirement categories. Preliminary information requirement categories are used in turn and cooperatively to formulate more specific information requirement categories which enable users and staff to establish just which journal, for example, will satisfy the information requirement. Again OPAC (and occasionally other resources to hand such as reading lists, notes and the contents of shelves) may be used to specify more precisely the nature of the information requirement. Over the course of this work, preliminary information requirement categories are often respecified – the information need may change from one located in a journal to one located in a reference book, for example – and the search area is further narrowed down.

4. Appealing to immediate user search history to perform categorization work in troublesome situations. Searches with help desk staff do not always run smoothly. Particular troubles arise in formulating specific information requirement categories. In such cases help staff desk routinely appeals to the user's immediate search history (what they have been looking at prior to turning to the help desk for assistance). In spelling out in detail just where users have been and just what they have looked at, their search history is employed by help desk staff both to eliminate areas of the search and to furnish new resources with which to elaborate and refine the search. Search histories are appealed to and elicited as a matter of course in the accomplishment of filtering work, providing for the narrowing down of the search through the cooperative formulation of specific information requirement categories and location of potentially suitable search materials.

Figure 3.9 The 'solution'.

The Pattern of Technology Usage

The online catalogue (OPAC) is used to work up information requirement categories and locate items that may satisfy the user's requirements. OPAC is a text-based system that allows vague descriptions of the information requirement provided by the user to be categorized in terms that fit the library catalogue. Vague descriptions are made to fit the catalogue through a number of static categories ('author', 'title', 'journal', etc.). OPAC provides lists of bibliographic items in response to queries issued (e.g. 'marketing journals', or 'java titles', for example), which users work through and employ to generate more specific information requirement categories and so refine the search. Notably, OPAC lists are randomly ordered and collect together all items containing the search terms employed (e.g. 'marketing journals', or 'java titles'). Thus users and help desk staff alike are presented with a vast array of topics which they must browse through in order to identify materials relevant to the current search, which takes a great deal of work.

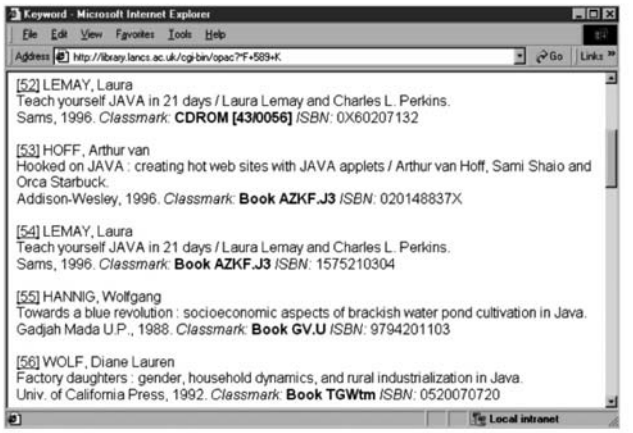

Figure 3.10 The 'diagram'.

action as elaborated by the work of the sequence and the practised ways in which that work is routinely organized.

The pattern section is moved up and placed after the *organizational context of the pattern*. This is done for reasons of relevance, as it is assumed that systems designers will primarily be interested in the pattern of technology usage. The other sections may subsequently be read and analysed for their relevance to the pattern. The 'tying paragraph' is renamed *connected patterns* and uses hyperlinks to connect the pattern

OPAC Use

[hyperlink to corpus of patterns of OPAC use listed in patterns index]

Figure 3.11 The 'tying paragraph'.

to a *patterns index*, which provides access to any other patterns in the setting that use the same key technologies and elaborates a bricolage of patterns of action that coalesce around particular technologies, thereby articulating important sites for design (Figure 3.11).

Unlike the use of pattern languages in software engineering or HCI, the adapted format offers no solution to designers (i.e., it doesn't tell designers what to build). This is quite intentional for as Vlissides (1997) reminds us,

> Clearly, one pattern format does not fit all. What does fit all is the general concept of pattern as a vehicle for capturing and conveying expertise, whatever the field.

The structure of the adapted format is relevant to the problem at hand, namely the requirements problem. No solution is offered by the format as it would conflate the design task; before devising technological solutions, we first need to establish a concrete sense of *what* is to be built. This means that we need to develop an intimate familiarity with the work of the site, with what goes on there and how it is organized. The adapted format is configured to address that question and so provides shared resources for analysing the design space and reasoning about the potential role of technology given the observable collaborative work of the site.

The adapted patterns format *mediates* communication between ethnographers and designers, providing a common focus for cooperative analysis of the design space. Patterns highlight important arrangements of cooperative work that the have been identified in ethnographic studies of the workplace and present these to designers in an accessible and relatively concise way (relative in respect of the thick ethnographic texts that are traditionally produced by ethnographers). In a similar manner to process maps (Crabtree et al. 2001a), patterns also serve to create a shared inter-subjective sense of the design space among the parties to design. The adapted framework supports analysis of the design space across the design team by illuminating the ecologically situated uses of technology and typificatory structures or common patterns of action that technology use is embedded in. Thus, the framework enables the members of the design team to identify the local constellations of cooperation and assistance shaping work in the real world and to reason about the demands that will be placed on design solutions. Patterns elaborate the actors, the concrete arrangements of collaboration they work within, and the sequential orders of cooperative work they routinely perform and so provide a solid basis to reason about user needs and to formulate use-cases or use-scenarios outlining potential design solutions supporting members' cooperative work (Crabtree et al. 2002a).

3.2.3 Analysing the Design Space with Patterns (Formulating Design Solutions 1)

Twidale et al. (1993) note that although various structured methods and CASE tools have emerged in recent years in support of software engineering, such approaches

> have provided very little support for the creative process of design, as distinct from solution structuring, refinement and documentation.

By the creative process of design Twidale et al. refer to analysis of the design space and the initial formulation of potential design solutions. Naturally, the creative process of design may be an ongoing process, one that is engaged iteratively as systems develop over time. Nonetheless, and as Twidale et al. point out, the creative process of design does not proceed through the use of formal methods in the first instance, but is carried out in "informal ways", the outcomes of which are subsequently "normalised" or made publicly intelligible and available to others through the use of formal methods. Rogers and Bellotti (1997) provide some instruction as to what that informal work consists of with respect to ethnographic involvement in the design process in saying that

> ethnography is most likely to show its value in being expounded within an ongoing dialogue between collaborating ethnographers and designers about observations and understandings derived from field studies, together with interesting capabilities of new technology configurations.

The cooperative analysis of patterns, upon which the initial formulation of design solutions stands, is not situated in the use of formal design methods, then, but in the *ongoing dialogue* that takes place between work analysts and systems developers in design meetings. The work of a perspicuous setting (a design meeting where support for search activities in the library is being considered) is consulted below in order to explicate what the work of analysis and formulation of design solutions observably and reportably consists of.

Library Design Meeting #1. (Edited fieldnote extract)

Jonathan: For me, every single search you perform on the library produces a new world, OK? You take that as my base thought. So you search by classmark, you search by title, you search by keyword, OK? Basically there are objects that come back, and they are labelled with their titles or whatever, and what I imagine is that you lay them out in a 3D space. You might bring the ones that are more commonly used closer to you. (Jonathan sketches a diagram of the 3D space on the whiteboard.) These are atomic objects, these aren't worlds, these are single objects, single books. (Points to objects in the diagram.)

Andy: I think that's something that's useful afterwards but there's a basic problem here. There's users who know exactly what they want, and there's users who have some sense of what they want but they don't know exactly as an item what they want. So what I know as a user is, I want something on Dada, for example, and I want something on the history of Dada but I don't know exactly what I want on the history of dada, but I want that. So you've done this search. Rather than getting two hundred and thirty objects back in various relations (on an OPAC list), what you are getting back is maybe five or six categories, and one of those categories is going to be history. That's what I want – history. I don't want philosophy, you know what I mean? What you are not having to do is wade through all possible things there (Points to diagram on the whiteboard.), and all those possible things consisting of irrelevant topics I'm just not interested in. This allows me at a glance to say, history, and go straight into that world.

Jonathan: Some of what you are saying is the common shortcomings of Q-pits. Of just – literally, you have a query and then you have a whole information space consisting of the results of that query, and then they are laid out *x, y, z* according to some criteria. The big problem with that is if you have got a billion objects it's meaningless. It's very difficult to get through it and hence, that's to do with what John and I have been talking about, about clouds – so you don't ever see, you don't get back a million objects, you get back several – four or five objects. They are labelled clouds based upon some commonality.

As the talk makes available, 'analysis' is a gloss on a lively context of work in which work analysts and designers formulate, discuss, argue over, contest, criticize, recognize problems with, reformulate, refine and otherwise negotiate the practical character of potential design solutions. Patterns of cooperative work are not simply relayed and responded to, then, but invoked and made relevant to design conceptions in the unfolding flow of talk. That work is quite observably done through *story-telling*. Story-telling may sound like a trivial way to talk about something as sophisticated as analysis in design. Nonetheless, design talk of all sorts, including the articulation of typificatory structures, is witnessably embodied in story-telling. Story-telling should not be treated lightly. Indeed, and as Erickson (1996) points out, the telling of stories is central to analysis of the design space:

> Stories work well as a way for promoting collaborative work and understanding within the design team. Stories are a sort of equaliser. It doesn't require much expertise or training to listen to and tell stories. Team members from any background can be made part of the process of telling and collecting stories. And once stories have been gathered, team members can discuss stories, argue about their interpretation, and [so on]. This process sensitises everyone to the usage domain, helps people identify questions and issues to probe ... and best of all provides concrete examples which can prove invaluable when

> team discussions threaten to veer into debates about vaguely defined abstractions.

Story-telling is central to the analysis of the design space and the formulation of design solutions because the telling of stories is *an ordinary method of analysis* (Sacks 1992b), which design work trades upon whether working with ethnographers or not. As an ordinary method of analysis, story-telling allows people with very *different* competences, expertise, knowledge, and professional languages – such as ethnographers and designers – to analyse situations together such that mutually intelligible understandings of the design space and user needs may be arrived at. The method permits cooperative analysis of the design space and diagnosis of problem situations. In and through telling stories, the ethnographer and the designer (and others of course) are able to analyse and construct a mutual understanding of the design space and user needs and to diagnose problem situations which technical solutions might be designed to support. Story-telling thus permits patterns of cooperative work to enter design reasoning in a readily digestible way. It does so, as the late Harvey Sacks (1992b) put it, because

> on analysing a situation they're in, [people] discover that they know about it with some story, which can be made, then, [into] something they ... have at hand, and may tell as a proposed analysis of [the] current situation.

So, for example, the situation that is articulated by the designer's story regarding how searching might be supported (through a system that allows keywords to be issued, that returns atomic objects, labels them according to various criteria, lays them out in 3D space, and places the most commonly used ones up front) is known about and analysed through the ethnographer's story that contests the proposed solution in telling of patterns of cooperative work and inherent problems encountered in searching (that users don't always know what they want, that OPAC lists are randomly organized, that users have to wade through or browse lots of irrelevant items to find what they want). The ethnographer's story also tells of a potential solution in light of those patterns and problems of work (that the 3D space might be organized in terms of an intermediary layer of topical categories that group atomic objects together). This story in turn becomes the focus of analysis and a potential solution is offered in the terms of Q-pit technologies that group objects together in clouds according to some criteria of commonality.

As the example indicates, the work of doing analysis through story-telling does not just consist in the telling and analysing of discrete stories but over the course of telling and analysing a series of stories, in the working up, elaboration and refinement of potential design solutions. Thus, it is through cooperative analysis and diagnosis of patterns of

cooperative work that work studies may come to give form to design in the first instance.[21] Story-telling is anything but trivial; indeed, as Orr (1996) points out, it is a much neglected area of research in occupational studies of all kinds. In the context of systems design, stories are unnoticed technical objects in the accomplishment of the day's work. Studies of their collaborative production and use promise to reveal much of design as a cooperative endeavour. That is not all there is to the bricoleur's craft, however, for although design solutions may come to take a preliminary shape through story-telling, that solution has yet to be fleshed out in detail.

3.2.4 Co-constructing Use-scenarios (Formulating Design Solutions 2)

Initial design solutions formulated in the course of considering patterns may be fleshed out in detail through the co-construction of use-scenarios (Carroll 1995). The production of scenarios is today a commonplace approach to working up design solutions. Just what the notion means, however, cannot be reduced to a simple set of work practices. A multiplicity of different forms of scenario exist – such as current work scenarios, future work scenarios, illustrative technical scenarios, evaluative scenarios, and so on – each of which plays a different role in the creative process of design (Bødker et al. 1995). Insofar as scenarios are an issue here, then it is largely with respect to future work scenarios, that is, to just what work will require technical support in the future and should, as such, be taken into account in the design of the new system (categorization work in supporting searching in the library, for example). Kyng (1995) describes such scenarios as use-scenarios. Whatever one chooses to call them, such scenarios are based on studies of work and

> indicate how computer support and (or) changes in work organization may improve upon work situations. Where work situation descriptions and overviews are about existing situations, use scenarios describe future possibilities ... these scenarios presuppose certain qualities of the computer support, and in this sense they also give some indication of requirements and thus resemble the use

[21]The working up of potential design solutions through story-telling is, more often than not, accompanied by the drawing of sketches and diagrams on pieces of paper or whiteboards, etc. Such artefacts tend to be very rough and informal and should be understood as temporally situated resources produced and used in the production of mutual understandings of the problem situation and of the general character of potential solutions, rather than as firm specifications for design. In that respect, such artefacts are subject to revision or amendment at a later date, or they may even be thrown away having served their purpose (positive or negative) in the temporal flow of design work. What the developers are left with (what is more enduring) are the 'war stories' produced over the course of design work, as these articulate in the re-telling, and in both a readily digestible and economical way, the rationale for what is being done and in the ways that it is (Orr 1996).

> cases of Jacobson. But our use scenarios are not fixed requirements; rather they represent hypotheses to be evaluated through [envisionment] workshops [with end-users] – and further revised and developed. Our use scenarios don't make sense in themselves ... [but] are understood through the mockups or prototypes realizing them. Thus, we don't develop our use scenarios *before* our mockups and prototypes, but concurrent with them. (Kyng 1995)

Use-scenarios might otherwise be described as open-ended use-cases. They do not fix design solutions but elaborate them over the course of their production, retrospectively through observations of work and prospectively through consultations with users. Just how use-cases may be worked up through consulting users will be addressed in the following chapter. First however, we need to appreciate what it means to base use-scenarios on studies of work, thereby grounding the formulation of design solutions in sequential orders or patterns of cooperative work.

By way of articulating what it means to base use-scenarios on studies of work, Kyng (1995) argues quite rightly that the role of proxy user suggested by some involved with ethnographic work is not tenable, but that does not mean that there is no room for ethnography in design:

> Doing initial analysis, developing work situation descriptions and use scenarios, preparing mockups for specific envisionment, and conducting workshops requires specific skills that are difficult and time-consuming to develop. Often designers, and obviously end-users, lack such skills. Thus cooperative design definitely needs facilitators, people skilled in working with different kinds of design representations and involving both users and designers in cooperative activities; people who can convince designers, end-users, and managers of the benefits of working cooperatively; in particular, people who can organize the workshops where mockups and prototypes of new design are tried out by end-users in work-like settings.

Kyng respecifies the perceived role of the ethnographer in the creative process of design from proxy to facilitator (by contrast, it might be argued that this has in various ways been the ethnographer's role all along but it is of little consequence here). Just what this means is rather vague, however, and one might say that just what it means concretely will be the outcome of doing the work, of engaging in design. In that respect it may be argued that the ethnographer's role as facilitator – as bricoleur – is not one of doing work studies and writing use-scenarios in light of those studies. Rather it is one of doing work studies and collaborating with designers in the joint formulation of use-scenarios. Many ethnographers are not in a position to write use-scenarios on their own or in collaboration with other work analysts. The ethnographer simply does not posses the technical competence needed to envision the future in light of the present. There is a necessity, then, to collaborate with designers in the production of use-scenarios.

Given the necessity for collaboration, it makes sense to speak of a division of labour in the production of use-scenarios. On one side of the division stands the designer, who brings knowledge of technology and technological possibilities to the design exercise. On the other side of the division stands the ethnographer, who brings knowledge of the social organization of work to the exercise (and, of course, users and other stakeholders need not be excluded from this work). Knowledge of these respective areas is originally articulated and combined through story-telling and the stories produced provide a basis for the formulation of use-scenarios, filling out in detail the practical character of design solutions (i.e., what the design solution will do and how). As a result of discussing proposed design solutions in relation to the sociality of work in the library, for example, the ethnographer and designer produced a story that served as a basis for developing a specific use-scenario. That story specifies systems support for searching's work such that on issuing a keyword search, items are returned, labelled, grouped, and linked together in navigable clusters of clouds according to the search criteria specified when the search is issued. This arrangement provides an 'intermediary layer' mediating between users and the catalogue, and furnishes material resources which may be employed by the parties to searching's work to accomplish categorization work.

Just how any of this is to be achieved technically is not the ethnographer's problem, not in the production of use-scenarios at least. The ethnographer's problem is not, as such, one of specifying how one should get from A to B as it were, but of specifying what A and B (and C and D, etc.) are in the first place. It is the designer's responsibility to make the technology answerable to these specifications. Thus, and for example, in considering technical support for searching's work the designer is required to articulate potential design solutions that will support categorization work. The extract below furnishes a concrete example of the use-scenario formulated by the ethnographer and designer with regard to supporting searching's work in the library.

Extract from Library Use-scenario

On executing a keyword search, the OPAC system will return a potentially large number of matches or 'hits'. The prototype is configured to group hits in the information view according to their classmark (which the system should be able to expand into more meaningful identifiers). Common groups of data (those with the same classmark) are represented as solid 'clouds' whenever the number of matches reaches a certain threshold (which is possibly user configurable and may also depend on system performance). In this example the threshold is passed and the system starts creating various clouds for 'Dada', such as 'history', 'art' and 'philosophy' for example.

The categorization used to form the clouds is user configurable (not necessarily by classmark). Thus a user may select some general attributes of the resulting hits (like classmark or author, etc.) or some other available metrics may be employed, like popularity (how many times have certain books been accessed before). In Figure 3.12 the clouds are defined using converted classmark-topic categories forming several areas, including art, philosophy and history – all of which contain books on Dada. The user chooses one of the clouds, history say, which contains a number of books within the cloud threshold value, and flies into it. The system presents the books (or objects) in the space (shown in Figure 3.13), and arranges them in three dimensions based upon some appropropriate Q-pit like criteria. For example, the popularity of use, insofar as the object has been marked as (potentially) suitable by users.

Once within the cloud, the user flies into several of the books, identifying two potentially suitable matches for their Dada search, and marks them for later physical retrieval from the library itself. These marked books are added to a session bookmark or shopping list which can be printed out by the user at the end of the session.

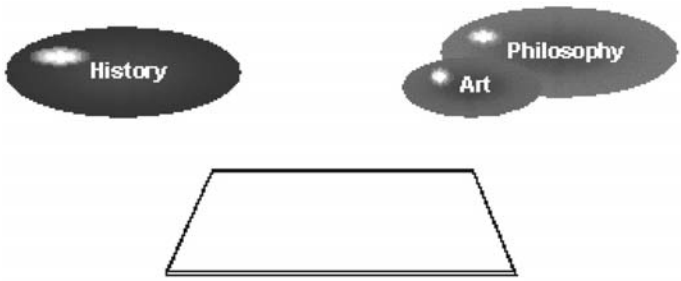

Figure 3.12 Clouding by classmark conversion.

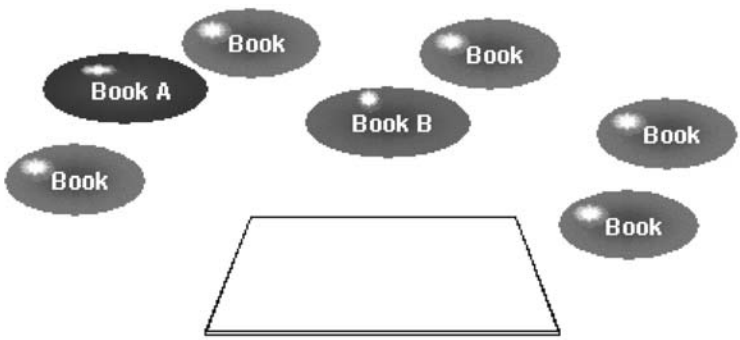

Figure 3.13 Presenting objects in the cloud.

It should be remembered that this is only a snippet from the overall use-scenario. Nonetheless, the extract hopefully makes it apparent how the requirements problem may be addressed through the organized use of work studies. Cooperative analysis of patterns of cooperative work in the library resulted in the intermediary layer story, which has been worked up in the use-scenario to lay out a *sequence of actions* that design solutions be may be modelled to support.[22]

As noted above, use-scenarios are developed hand in hand with mock-ups and/or prototypes. Describing sequences of action to be supported by technological solutions, use-scenarios provide concrete resources for configuring working prototypes. This configuring work draws contingently upon a host of different competences to hand, such as knowledge of legacy systems, database configuration, object-oriented design methods, and more, all of which are directed towards implementing technical solutions supporting the sequences of action articulated in the use-scenarios. The configuration of mock-ups and working prototypes is often conducted at whiteboards and with ethnographers and other relevant parties in order to ensure that technical solutions address the problem situation. Notably, it is in the course of this work that formal designs are sketched out (as portrayed in Figure 3.14, for example). With mutual understandings of the problem situation and formal design sketches in hand, the designers may deploy formal methods to implement the solution, employing CASE tools, graphical notations, and whatever else is necessary to configure a working prototype.[23]

The work involved in using formal methods is not considered here. Attention has instead been paid to the 'seen but unnoticed' or taken for granted and ignored work involved in design, particularly the appeal to typificatory structures (i.e., patterns of action), the nature and role of story-telling, and its relation to the co-construction of use-scenarios, all of which presuppose the deployment of formal methods. With use-scenario(s) in hand, a suitable architecture may be sketched out and for-

[22]The co-construction of use-scenarios is an iterative exercise in which documents are worked up, elaborated and refined by the ethnographer and designers over a series of turns in an ongoing course of dialogue.

[23]Formal design work at whiteboards in design meetings is one area of cooperative work in design not considered in detail here. It is an important area of research in its own right and, indeed, has received attention elsewhere in the design literature (e.g. Pedersen et al. 1993). Its description and analysis require a vulgar competence in formal design methods (which the author does not possess, hence the absence of a detailed account). Nonetheless, it is at this site that the 'formal' and 'informal' often meet. The ethnographer need not be unduly concerned by this work, however, as the implications of the work studies have already been articulated in an intelligible and digestible fashion to designers in the use-scenarios. As noted above, the work here is, amongst other things, concerned with meeting the constraints of the use-scenarios and the ethnographer's role is an auxiliary one of clarifying the use-situation as and when particular problems emerge in the design dialogue.

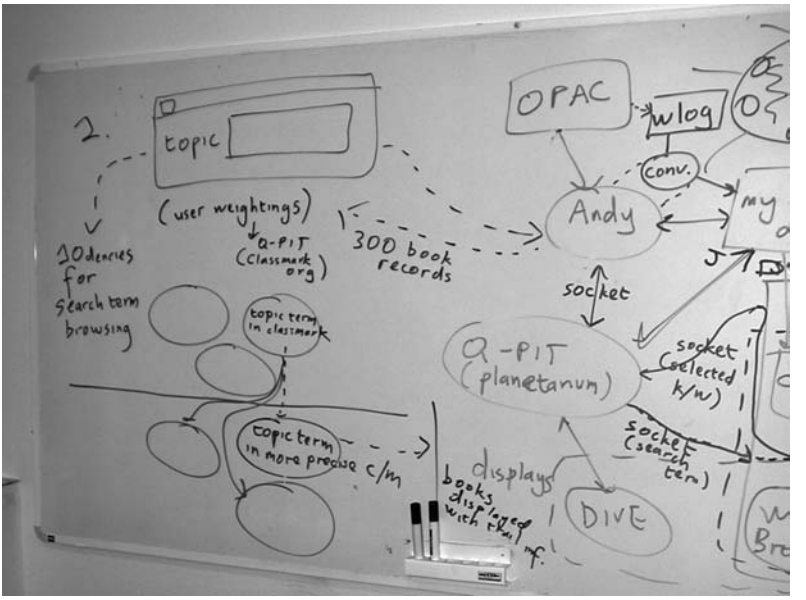

Figure 3.14 Configuring scenario-based prototypes.

mal methods deployed in the design of a technology that may be used to support, in this case, searching's work. In this case, the proposed design solution consisted of developing a 3D visualization that reflected the core elements of the use-scenario, and which interfaced to the library's catalogue system. This configuration allowed users to access real world data through the 3D environment and to use the visualization of data structures to accomplish search activities. The core elements of the working prototype are presented below prior to considering the role of end-user evaluation in the design of collaborative systems.

The working prototype exploits Q-pit technologies (Mariani et al. 1995) running in DIVE (Carlsson and Hagsand 1993) and DEVA (Pettifer et al. 2000) environments. These technologies allow multiple users to simultaneously inhabit a shared 3D environment and have, as such, tended to be seen as social spaces where the purpose is to allow users to meet each other and interact together within 'cyberspace'. The notion that collaboration and cooperation should take place *within* the virtual environment is a prevalent one, guiding a great deal of research in the design of Collaborative Virtual Environments (CVEs). However, studies of CVE usage point out that workaday activities are achieved by a combination of collaborative actions taking place in the virtual environment and in the real world (Bowers et al. 1996). In light of this it is suggested that

> when it is claimed that CVEs can in principle support cooperative work in ways difficult to achieve with alternative technologies, we take this as a claim most appropriately assessed in light of *all* the work *both within and without* the virtual environment narrowly considered. CVEs should not be criticised (nor prematurely celebrated) on the basis of *only* what can be designed into the virtual environments or occurs within them. This argument, then, suggests that there is an opportunity to rethink the design challenges in CVE research: *one should be designing for two worlds not just one*. (Bowers et al. 1996)

Comments such as these suggested that the design of CVEs should be extended to provide support for the cooperative work that takes place outside the virtual environment in the real world. In attending to the external world of work, more recent developments have seen emphasis placed on *augmenting* the material resources in and through the use of which cooperative work gets done (resources such as OPAC, for example). As Büscher et al. (2000) put it,

> CVE researchers have experimented with different ways of mediating awareness and formation of groups and meetings in virtual environments. However, the artefacts related to the work of users have received less attention ... This is a serious barrier for bringing CVEs out of entertainment or research to real world work tasks.

Taking the material demands of cooperative work in the real world seriously suggests that the major challenge in the development of CVEs is not so much to focus on what can be done cooperatively within the environment, but to support cooperation in the real world through the augmentation of material resources. The point was taken seriously in configuring the library prototype, and emphasis was accordingly placed not on what should or could be done cooperatively within the virtual environment, but on providing support for real world patterns of cooperative work. Putting novel technologies to work by *situating* them within real world patterns of cooperative work was taken to be the major design challenge and the working prototype was therefore configured to support work as it is, rather than a priori concepts of what CVE technologies should do. In that respect, it was noted that the library catalogue is accessed through public terminals which people cluster around in various constellations of cooperation and assistance. This shoulder-to-shoulder interaction is central to how OPAC users conduct their work together and the prototype was configured to be shared in this shoulder-to-shoulder manner rather than remotely and through the use of avatars.

In contrast to OPAC, which provides a number of fixed entry categories to input a search term, users issue searches via the developed prototype through a single keyword entry window (Figure 3.15). This reflects the topic-based character of searching and provides the basis for the generation of the virtual environment. In order to issue a search the user simply

fills in the keyword slot. Items are returned according to default search criteria or 'weighting'. The user may change this weighting so that items are grouped by similarity according to author, title, publication type, or date, publisher, ISBN, or any combination of these criteria and by any other criteria that may be added to bibliographic items in the future (e.g. via HTML Metatags). The user may also adjust the degree of exactness by which similarity weightings are applied (from a range of high to low) and the default threshold which determines the numeric level at which grouping takes place such that if only ten items are returned, say, then they are simply presented to the user as atomic objects, whereas if seventy items are returned, then they are grouped according to the similarity criteria specified.

Once a keyword has been provided, the search may be executed within the OPAC catalogue. The results are then used to populate the environment. The final stage of the query and retrieval process consists of the construction of the intermediary layer suggested in the design meeting through formation of links and clouds in the space. The clouds display at-

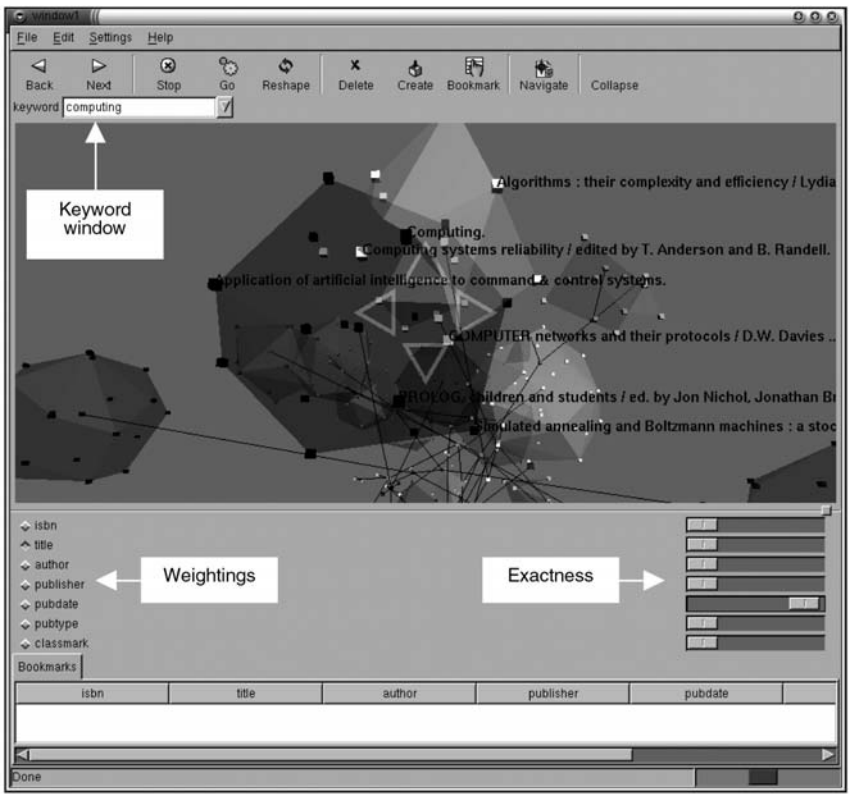

Figure 3.15 The Q-pit library interface. © 2000 Steve Petifer (AIG, The University of Manchester). Included here by permission.

a-glance groupings of similar items (similar in respect of the weightings applied to the search) and the links indicate interconnections between clouds and items therein. Each object within the space has a partial title attached to it. As clouds of similar objects form, users can view object titles and see at a glance just what each cloud contains as a topical space. In issuing a keyword search on 'java', for example, the user can see by viewing the titles of the objects being pulled into the clouds that one of the clouds contains items on Java software engineering. This may contrast with the objects in an interconnected cloud whose titles indicate that objects on Java software testing are contained within it. As the search engine pulls all items with a similar keyword into the space, objects about Java the island may be pulled into the space as well and would be pulled together in a separate, unlinked cloud, although that cloud may be linked to other Java-the-island topics. The grouping of related items contrasts with the OPAC interface, which groups everything together in a randomly organized list.

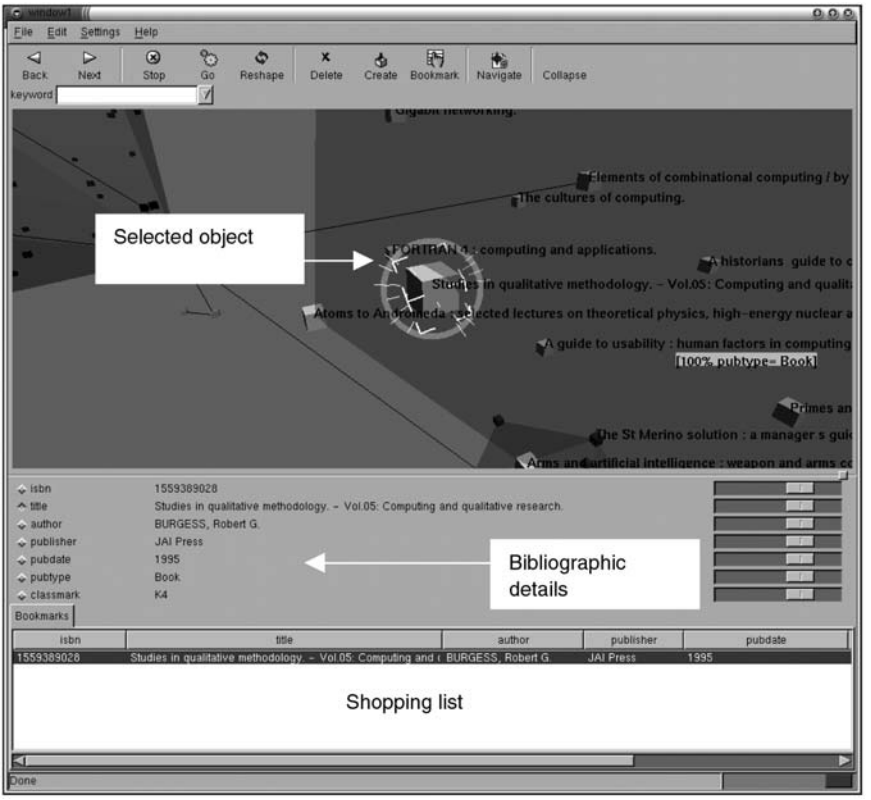

Figure 3.16 The Q-pit library interface. © 2000 Steve Petifer (AIG, The University of Manchester). Included here by permission.

Once the space has been generated, users are free to interact with and navigate the display. A key aspect of the environment is that users have an object-centric view of the space. In other words, navigating or searching the space is constrained in terms of specific objects and the relationships that hold between them. In contrast to other virtual environments where full 3D navigation is provided and often found problematic, the real world nature of the application suggested the need to support a constrained form of navigation in order to avoid the overhead of managing movement in a virtual environment, which is notoriously difficult to support; users are constantly getting lost in unconstrained environments, for example. Consequently, an object-based approach to searching the space was exploited with explorations centring on a selected object (Figure 3.16). In this respect, the navigation vehicle developed supports two modes of interaction: 1) browse, and 2) point and click. In the browse mode, mouse movements allow the user to orbit the clouds in the space and see the various constellations of topics and their interconnections. Users may examine particular objects by centring the object in the cross-hairs to see bibliographic details. Over the course of interacting with the space, the user can also assemble a shopping list of items that are of potential relevance or interest. The list may be printed off and used in the course of locating objects in the physical library. The shopping list persists across different searches and is, as such, session-oriented, providing support for categorization work at the help desk in times of trouble, which trades on the user's search history. Each search space may be reordered or regenerated by simply altering the weightings to display different relations of similarity between the displayed search items. From this point, regeneration continues much as for the original generation of the space (Crabtree et al. 2002b).

Developing a scenario-based prototype is by no means all there is to the formulation of design solutions. In many respects, this might be said to be just the beginning of the enterprise, as the developed environment articulated above is but the first incarnation or version of the design solution supporting searching's work. The point of describing the work involved in the production of this early and incomplete prototype has been to illustrate how ethnography might support analysis of the design space and shape the formulation of design solutions. The production of patterns, working up of stories, and the co-construction of use-scenarios illustrate practical ways in which ethnography may be deployed to inform design reasoning and the production of design solutions.

In a design context, the disciplinary concerns of social science are left behind, which is likely to be a great disappointment to social scientists who see design as a promising new site in which to pursue their agendas. In that respect, the work articulated here should serve as a caveat to designers and social scientists alike. Designers have unique problems of their own and there is no need to suppose that they are identical to those

that beset the agendas at work in the social sciences. On the contrary, design is not interested in social science but in developing methods for analysing the sociality of work in the wild, in its natural habitat, in contrast to in social science theory. This chapter has articulated several ways in which the development effort may be constructively informed through ethnographic studies of work. This is not to say that the approach advocated here is the only way in which the social sciences in general, and ethnography in particular, may enter design reasoning – it is but one way and one that complements existing design practices which presuppose the deployment of formal methods and rely on ordinary competences (and so require no long-term specialized training). The purpose of the next chapter is to consider what more there is to the ethnographer's role as a bricoleur in design practice by addressing what more there is to the formulation of design solutions supporting cooperative work.

An alternative approach to traditional requirements specification relies on ethnographic studies and/or user participation throughout the design process to minimize the mismatch between the computer system and conditions of use. How and in what ways ethnography and techniques of user participation can be effectively developed is a major research area within CSCW.

Liam Bannon & John Hughes