

On “Technomethodology”: Foundational Relationships Between Ethnomethodology and System Design

Paul Dourish

Xerox Palo Alto Research Center

Graham Button

Xerox Research Centre Europe

ABSTRACT

Over the past 10 years, the use of sociological methods and sociological reasoning have become more prominent in the analysis and design of interactive systems. For a variety of reasons, one form of sociological inquiry—ethnomethodology—has become something of a favored approach. Our goal in this article is to investigate the consequences of approaching system design from the ethnomethodological perspective. In particular, we are concerned with how ethnomethodology can take a foundational place in the very notion of system design, rather than simply being employed as a resource in aspects of the process, such as requirements elicitation and specification.

Paul Dourish is a computer scientist investigating the design of collaborative technologies on the basis of sociological investigations of work; he is a Member of Research Staff in the Computer Science Laboratory at Xerox PARC. **Graham Button** is a sociologist whose research interests are framed by ethnomethodology and conversation analysis; he is a Principal Scientist and Area Manager for the Studies of Technology, Organisations and Work group at Xerox Research Center Europe, Cambridge Laboratory.

CONTENTS

- 1. INTRODUCTION**
 - 2. WHAT IS ETHNOMETHODOLOGY?**
 - 2.1. The Origins of Ethnomethodology
 - 2.2. The Objective Reality of Social Facts
 - 2.3. The Accountability of Social Action
 - 2.4. Membership
 - 2.5. Ethnomethodology and Ethnography
 - 3. THE RISE OF ETHNOMETHODOLOGY IN HCI RESEARCH**
 - 3.1. Plans and Situated Actions
 - 3.2. Participatory Design
 - 3.3. Ethnography
 - 3.4. HCI in Transition
 - 4. ETHNOMETHODOLOGICAL STUDIES OF TECHNOLOGICAL WORK**
 - 4.1. Learning From Ethnomethodologists
 - 4.2. Learning From Ethnomethodological Accounts
 - 5. ETHNOMETHODOLOGY FOR CRITIQUE AND DESIGN**
 - 5.1. Two Paradoxes of Ethnomethodologically Informed Design
 - 5.2. Critique and Design
 - 6. TECHNOMETHODOLOGY: DRAWING FOUNDATIONAL RELATIONSHIPS**
 - 7. ABSTRACTION, ACCOUNTS, AND ACCOUNTABILITY**
 - 7.1. Abstraction in System and Interface Design
 - 7.2. Abstraction and Glossing Practices in Everyday Interaction
 - 7.3. Toward Observable-Reportable Abstractions: Open Implementation
 - 7.4. Accounts and Accountability
 - 8. EXAMPLE: ACCOUNTING FOR FILE COPYING**
 - 8.1. The File-Copying Scenario
 - 8.2. An Account of File Copying
 - 8.3. Accounts and Ordinary Operation
 - 8.4. Accounts and Mental Models
 - 8.5. Accounts as an Example of Technomethodology
 - 9. ABSTRACTION AND PARTICULARITY**
 - 10. CONCLUSIONS**
-

We begin by outlining the basic elements of ethnomethodology and discussing the place that it has come to occupy in computer-supported cooperative work and, increasingly, in human-computer interaction. We discuss current approaches to the use of ethnomethodology in systems design, and we point to the contrast between the use of ethnomethodology for critique and for design. Currently, understandings of how to use ethnomethodology as a primary aspect of system design are lacking. We outline a new approach and present an extended

example of its use. This approach takes as its starting point a relationship between ethnomethodology and system design that is a foundational, theoretical matter rather than simply one of design practice and process. From this foundation, we believe, emerges a new model of interaction with computer systems, which is based on ethnomethodological perspectives on everyday human social action.

1. INTRODUCTION

One of the more significant trends in human-computer interaction (HCI) research and practice over the past decade has been the increasing influence of sociological perspectives in the design and evaluation of interactive systems. Sociological understandings and methods have been used to study the settings in which work is conducted, to inspire and guide the design of interactive systems, and to evaluate those systems in real working conditions. The uptake of sociological research has been most pronounced within the domain of Computer-Supported Cooperative Work (CSCW), where, clearly, sociological approaches lend themselves well to a primary focus on interaction between individuals and groups, rather than simply between the "human" and "computer" of HCI. CSCW research has highlighted the social setting of computer use and so set up for itself a framework within which sociology has direct applications. More recently, however, sociological perspectives have begun to permeate HCI more generally. Within HCI, sociological methods complement (and sometimes challenge) the technical and psychological perspectives around which the field was originally organized and have become increasingly accepted and even expected as a component of HCI research.

It would be unwise to imagine, however, that "sociology" is all of a piece—far from it. There are any number of particular branches, each with their own perspectives, orientations, methods, and concerns, coming together under the sociological umbrella. Shapiro (1994) provided something of a "travel brochure" for some of this vast terrain. In this article, we are concerned with one particular branch of sociological investigation: ethnomethodology. For a variety of reasons, some of which we explore here, ethnomethodology has become something of a favored (or, at least, more prominent) perspective among the sociological positions exploited in CSCW and HCI.

We come at this as representatives and practitioners of each of the two disciplines under consideration here. One of us (Dourish) is a computer scientist, and one (Button) an ethnomethodologist. Over the past few years, in work conducted separately and together, we have been concerned with how the design of collaborative and interactive systems can be grounded in sociological understandings of action and interaction. Our goal has not simply been to develop a model of design that is responsive or respectful to observations of par-

ticular social settings, and we have not been attempting to formulate a design method by which sociologists and computer scientists can work together on design problems. Rather, we have sought to develop a form of technological design that is fundamentally grounded in the understandings that sociological perspectives employ. Our fundamental position is that the relation between social and computational sciences is more than a practical problem. Our goal has been to develop a stance in which ethnomethodology and computer science play equally significant roles (rather than grafting one onto the other), and so our approach is radically novel for both disciplines. We use the term *technomethodology* to emphasize that it is something new, drawing from each side, but different from each.

Our goal in this article is to motivate, introduce, and illustrate the approach at which we have arrived. As a consequence, this article is rather unusual among those in HCI or CSCW emerging from collaborative work by sociologists and computer scientists. We present neither a set of technical requirements derived from a field study nor a system design that incorporates the lessons of ethnographic investigation. Instead, we are concerned with the basis on which those other sorts of research can be conducted. Because we are more concerned with computational design per se than with specific system-design efforts, our primary illustration is not a particular system, but a reconceptualization of a particular foundational element of interactive system design (i.e., the notion of abstraction). Before we can proceed to this, however, we need to spend some time considering just what ethnomethodology is and how it has come to play its current role in HCI.

2. WHAT IS ETHNOMETHODOLOGY?

Despite the interest in ethnomethodological ideas in CSCW and increasingly in HCI research, those ideas themselves remain remarkably poorly understood. We can speculate about the reasons. Perhaps it is because ethnomethodology largely has been conducted by ethnomethodologists, rather than becoming a more available approach to researchers at large, or perhaps it is due to the relative opacity of much of its writings: As Livingstone (1988) commented, "I know of no discipline which has suffered more at the hands of its expositors than ethnomethodology." In this section, we attempt to give a flavor of the ethnomethodological position,¹ with two concerns: first, to introduce some conceptual foundations on which we later build, and second,

1. Readers who wish to pursue this in more detail are referred to Garfinkel (1967) or Heritage (1984).

to give a flavor of the way in which ethnomethodology differs from other approaches.

2.1. The Origins of Ethnomethodology

Historically, ethnomethodology has its roots in the work of Harold Garfinkel, beginning in the late 1950s and subsequently developed through the 1960s and early 1970s by Garfinkel and colleagues, perhaps most notably Harvey Sacks, the founder of Conversation Analysis (CA). Garfinkel's objective was to respecify the subject matter and methodological approaches of sociology. At the time, the prevailing school of (especially American) sociological thought was structural-functionalism, most fully developed at that point in the work of Talcott Parsons, and especially his 1937 *The Structure of Social Action*. Ethnomethodology arises from Garfinkel's confrontation with this perspective: He later wrote, "Inspired by *The Structure of Social Action*, ethnomethodology undertook the task of respecifying the production and accountability of immortal, ordinary society" (Garfinkel, 1991, p. 11). So Garfinkel's project was not merely to critique Parsons, but to use Parsons's perspective as a starting point from which to question the very nature of what sociology was, what questions it addressed, and how it went about answering them.

2.2. The Objective Reality of Social Facts

A critical focus of this respecification is ethnomethodology's rejection of the traditional approach to the relationship between practical social action and the sociological "rules" by which stable social order is established and maintained (known as the "problem of social order").

Durkheim said, "The objective reality of social facts is sociology's fundamental principle," and from this principle, all sociological reasoning and sociological practice followed. Because social facts were, axiomatically, objectively real, sociology could go about studying what those facts were and how their consequences played out. It was this model of sociology that Parsons elaborated. For Parsons, the problem of social order was a matter of concerted action, and so he went about describing how, in performing activities in accordance with reciprocally shared rules and norms, social actors achieve the coordination of their activities.

Garfinkel's "respecification" struck not simply at the work that Parsons and his colleagues had done, but at the very foundations that they had drawn on—the axiomatic "objective reality of social facts." Garfinkel was not satisfied with the idea that stable social order proceeded naturally and uncomplicatedly from those social facts. He drew attention to the work that goes into the production of social order, underscoring how social order is "made to work" in the actions and

interactions of its members. Social order does not simply exist, and social action is not simply determined from it; and so, social order and social action cannot be approached independently.

So, for ethnomethodology, the “objective reality of social facts” is not sociology’s fundamental principle, because it is not a principle at all; it is sociology’s fundamental phenomenon. It is not to be assumed, but to be studied. Precisely how this phenomenon is achieved and manifested became ethnomethodology’s primary subject matter for investigation. So ethnomethodology turned its analytic attention to the ways in which everyday social action was achieved, looking directly within the circumstances of action for evidence of the methods by which individuals achieved precisely the stable social order that traditional sociology had defined by theoretical fiat.

2.3. The Accountability of Social Action

In looking at the emergence of social facts from the everyday details of what people do, a primary concern for the ethnomethodologists was not just how individuals engaged in rational social behavior, but also how they could be seen by other actors to be engaged in it. After all, the ethnomethodologists reasoned, “rational social behavior” was observable not only to learned professors of sociology who knew “the rules,” but also to individuals engaged in everyday practical action. Everybody, everyday, knows what rational social behavior is and what it looks like, even though they have never studied sociology; the question is, How?

The phrase “everyday practical action” is a telling one. By emphasizing everyday action, ethnomethodology draws attention to the fact that its concern—the production of social order—is no special activity. It is not done only by certain groups (like sociology professors) or only at certain times (like revolutions); it is a part of ordinary, everyday life, woven into the fabric of all activity. Similarly, the emphasis on practical social action is both intentional and highly significant. It implies two things. First, it emphasizes that, as social action unfolds in people going about everyday activities, the activities are where their interests lie. In other words, there are matters at hand to be attended to, and the object of activity is to attend to them, not to reproduce the stable structures of society. When I call a plumber and ask him to fix a leaking pipe, my concern is to avoid a pool of water on the floor, not to reproduce a pattern of social interaction based on contracts and wage-labor. Second, focusing on practical social action draws attention to the way in which those practical concerns, the matters at hand, take up a critical role in the understanding and production of action. Understandings of social action are formed “for practical purposes” by the participants; they help us get the job done.

So, everyday practical action involves not just engaging in rational social behavior, but also being seen by others to be so engaged. The "rationality" of social behavior lies in the way that it is intelligible to others. Garfinkel held that the procedures for the production of action and the procedures for observing people as being engaged in rational social action were one and the same.

Ethnomethodologists capture this by saying that human social action is "reflexively accountable"—in other words, that, first, the very way in which it is organized provides to others the means to recognize it as what it is (accountable), and second, it does so within the very fact of its production (reflexivity) rather than within some wider frame of "social meaning". So, the organization of action, within specific (and innumerable) circumstances of its production, furnishes to others the means to recognize, observe, and report on it. For example, an "ordinary social fact," such as a conversational greeting, is produced in conversation as an observable greeting, not just by using a recognizable greeting term such as "Hello," but also by placing that term in an interactionally organized position within the course of a conversation. The "work" that the word does is a feature of how the word is used rather than simply what the word is. In other circumstances, "Hello" can be a request for attention, a howl of derision, an inquiry, a mark of interest, a signal of surprise, and so forth. The question is how it is used, and how it is heard as being used one way or another. Ethnomethodology observes that the circumstances in which language is used to perform social action grant to other participants the means to recognize the nature of that action. As a general concern, this property of human social action—accountability—became one of the critical analytic features of ethnomethodological studies.²

The idea of the reflexive accountability of action also provided ethnomethodology with an analytical warrant for a particular form of investigation, one which considered specific instances of action in extreme detail, looking within those circumstances for the exhibition of members' methods of acting and seeing. This is, perhaps, most vividly demonstrated in the development of CA, a particular form of ethnomethodology, developed primarily by Sacks (1992).³ CA studies take fragments of naturally occurring conversation as data,

2. In our experience, one confusion about ethnomethodology in HCI has arisen precisely from Garfinkel's use of the term *accountable*. As we have outlined here, the accountability of action lies in the way in which it "gives an account of itself as itself," and is "observable and reportable" as such. It does not refer, then, to a political or moral accountability for one's actions and does not provide for an opportunity to be taken to task over them. The fact that action is accountable has nothing to do with the fact that someone may be "held accountable" for it.

3. CA has been exploited in the design of interactive systems largely independently of the recent influence of ethnomethodological ideas (see, e.g., Frohlich & Luff, 1990).

looking within the conversational data itself for the mechanisms by which conversation was systematically organized. Reflecting ethnomethodology's concern with practical action, CA begins with a perspective on conversation as social action, rather than as the articulation of internal mental states; on that basis, it analyzes this action to see how aspects of conversation (such as introducing a new topic, or bringing a conversation to a close) are managed as a practical activity. So, for example, in a CA investigation of greetings, the focus is not on what sorts of terms (like the word *hello*) might carry "greeting-ness" as intrinsic properties that people then deploy in different circumstances. Instead, it focuses on the interactional work that specific utterances do, the implications they have for what comes next, and for how they are used in the solution of the problem of how to conduct concerted social action on this occasion.

Because accountability is a fundamental and "irremediable" property of social action, the ethnomethodologists contended that this kind of analysis could be conducted not only in the domain of conversation, but in any domain of social action. Their concern with how social action was made to be rational, and made to be seen to be rational, allowed them to treat any social action as an occasion of lay "sociological theorizing." Conversely (and perhaps rather mischievously), it also allowed them to treat "professional sociological theorizing"—that is, the professional practice of sociology, carried out in books, journals, conferences, and lecture halls—as just as another domain of everyday practical action. Although this might well have been an interesting example of the universality that the ethnomethodologists were seeking, it has not endeared them to other sociologists!

2.4. Membership

From their concern with conversation and with language as social action, Garfinkel and Sacks (1970) emphasized a focus on "member" not as an instance of a delineable social grouping but as "mastery of natural language" (p. 343). For them, the mastery lies not in the grasp of syntax or grammar, but rather in the competent use of language (that is, of course, social action) in appropriate (social) settings. Such competent language use means to be able to use language, or to act, appropriately for the setting, which in turn depends on the appropriate exercise of understandings of (in Garfinkel's term) "what everyone knows that everyone knows."

For example, in his essay, "On Doing 'Being Ordinary'" Sacks (1984) illustrated the way in which these sorts of understandings are applied in everyday conversation. Sacks explored fragments of a conversation between two friends discussing a police incident at a store and a car accident and showed how the conversation is formulated so as to render explicitly some elements of everyday life as commonplace and ordinary, whereas others are rendered particular

and exceptional. These sorts of practices unfold against (and themselves recreate) a background of commonly understood ordinariness shared by the speakers, and so they are firmly situated in the circumstances within which the conversation itself takes place.

More generally, the notion of "what everyone knows that everyone knows" speaks to a form of commonsense understanding ("common" both because it is shared and because it is mundane) that is the basis of the mutual recognizability of accountable action. Ethnomethodology focuses on how people exercise these commonsense understandings by finding, within the immediate circumstances of action, the means to understand them and interpret them for practical purposes. These different elements—accountability, membership, and commonsense understandings—together contribute to the ethnomethodological frame of reference that places it in opposition to the view of the work espoused by Parsons and traditional sociology.

We only had space here to give the briefest outline of the ethnomethodological perspective, focusing in particular on those aspects that uniquely define its oppositions to traditional perspectives on social action. Our goal has been to set out enough of the background to frame further discussion of the role that ethnomethodological ideas can play in the design of interactive systems and particularly the core element of the practical, situated, accountable, and ordinary character of social action.

2.5. Ethnomethodology and Ethnography

When considering the role that ethnomethodology has played in HCI, it is important to make clear the distinction between ethnomethodology and ethnography. The fact that ethnomethodologists often use ethnographically generated materials in their analysis (Hughes, Randall, & Shapiro, 1993) may lead those who are not sufficiently familiar with the disciplines to conclude that they are the same thing. They are decidedly not.

As outlined previously, ethnomethodology is a particular analytic orientation to the practical issue of the problem of social order. It sets out a policy for the study of practical social action. Ethnography, on the other hand, is a form of investigative fieldwork and analysis. Ethnography considerably predates ethnomethodology. Modern ethnography emerges primarily from the work of Bronislaw Malinowski in the early part of this century, particularly his work in the Trobriand Islands during World War I.⁴ Ethnography is best seen in

4. The historical context of Malinowski's work is a different but fascinating story in itself. Malinowski spent the war years in the Trobriand Islands to avoid internment as a foreign national at the outbreak of hostilities (see Anderson, 1996).

contrast with other methods in anthropology; qualitative rather than quantitative, with an emphasis on the "member's point of view," and, critically, with a focus on the member's experience rather than simply his or her action. Ethnography has grown to be the predominant perspective of anthropological fieldworkers, not simply for the collection of their materials, but also for their organization, interpretation, and presentation. Within ethnography, however, numerous analytic orientations may operate. For instance, ethnographic field techniques also have been used by many in the "Chicago School" of Human Ecologists and Symbolic Interaction in the study of social life and of work.

Part of the confusion between these terms and approaches in HCI rests on the fact that ethnomethodology often makes use of ethnographically gathered materials. An ethnomethodologist going into the field to collect data is likely to use ethnographic techniques, and so, to an observer of fieldworkers, might seem indistinguishable from, say, a symbolic interactionist doing the same. Indeed, they might both be adequately labeled "ethnographers."⁵ The point of difference comes into play in the "analytic mentality" they display in the selection of phenomena and topics for investigation and in the issues they would want to draw attention to in the materials gathered. Some of the confusion, then, arises in the way in which these concerns have entered the domain of HCI research, which is the topic that now concerns us.

3. THE RISE OF ETHNOMETHODOLOGY IN HCI RESEARCH

Ethnomethodology has become a prominent form of sociological analysis in HCI and CSCW. This is particularly intriguing for us because, as should be clear from the previous discussion, ethnomethodology is only one among a wide range of sociological perspectives (and a fairly small one at that). Our goal here is to consider why, other than a partisan belief that ethnomethodology is a methodologically and analytically superior form of sociological reasoning, this has come to be the case.

More or less, ethnomethodologically oriented investigations now are presented regularly at CSCW conferences⁶ and increasingly at HCI confer-

5. Care must be taken here, too. Anderson (1991) made the point that what system design often sees as the value of "ethnography" is often simply the value of fieldwork, and Anderson discussed how ethnography itself also comes with its own analytically and politically predispositional baggage.

6. Examples include Anderson, Button, and Sharrock (1993); Bentley et al. (1992); Bowers (1994); Bowers, Button, and Sharrock (1995); Grinter (1997); Heath, Jirotko, Luff, and Hindmarsh (1993); and Rouncefield, Hughes, Rodden, and Viller (1994).

ences.⁷ HCI and CSCW being design-oriented disciplines, ethnomethodology is being used to inform design through

1. Fieldwork investigations that develop an understanding of work and organizations from the "inside," providing innovative insights into the organizational situatedness of work and the methods and practices through which work activities and interactions are assembled and that may be used in the design of technology to support work.

2. The development of an understanding of the temporal organization of activities and interactions, revealing them to be a moment-by-moment organization and, in so doing, furnishing new concepts around which to generally consider the design of technology.

These understandings allow a new focus on the relationship between technology and the accomplishment of work, one that emphasizes the technology as a part of the circumstances of the production of working order. From this analysis comes the opportunity to use ethnomethodological analyses as the basis of design and redesign of interactive technologies.

There are a number of interrelated reasons why some within the design community are taking up these two sets of issues.

3.1. Plans and Situated Actions

One primary reason for the widespread influence of ethnomethodology in interactive systems design is the role of Lucy Suchman's (1987) book *Plans and Situated Actions*. The book formulated a telling and forceful critique of the user modeling and planning-based approaches common in both HCI and Artificial Intelligence. It is very widely read and cited in the HCI literature, and it firmly established the relevance of sociological and anthropological reasoning for the problem of HCI. As such, the book and the argument Suchman puts forward have come to occupy an almost iconic position within the field, one which, due to a number of misunderstandings, Suchman has repeatedly been forced to clarify in the decade since its publication.

The disturbingly common caricature of her position is that there are no plans, but only "situated actions"—improvised behaviors arising in response to the immediate circumstances in which actors find themselves and in which action is situated. In fact, as Suchman has been at pains to point out, she did, in

7. Examples include Bowers and Pycock (1994); Button and Sharrock (1995a); Heath and Luff (1991); Hughes, O'Brien, Rodden, Rouncefield, and Blythin (1997); and Sellen and Harper (1997).

fact, accord an important status to plans as resources for the conduct of work. Her argument was that plans are one of a range of resources that guide the moment-by-moment sequential organization of activity; they do not lay out a sequence of work that is then blindly interpreted.

The argument that Suchman presented in *Plans and Situated Actions* was partly founded in the work of ethnomethodology. Garfinkel (1967) used the term *judgemental dope* (p. 66) to characterize traditional sociology's view of members' practical decision making, as they blindly act in accordance with theoretically formalized systems of rules and norms. Garfinkel attempted to relocate practical decision making to a realm of relevantly invoked situated actions in local circumstances. Similarly, Suchman emphasized a perspective on purposeful human action as situated in (and organized around) the context of particular circumstances. Suchman illustrated her argument with detailed examples drawn from a laboratory study of the use of a complex photocopier, and, as she pointed out, laboratory studies are hardly the stuff of ethnomethodology. However, her argument and analysis drew strongly on the ethnomethodological tradition and introduced it to the HCI community. The HCI community has never recovered. In a recent book collecting essays on different social perspectives on HCI (Thomas, 1995), 11 out of the 12 essays cited *Plans and Situated Actions* or Suchman's subsequent work.

Suchman's book has had a significant influence in HCI design and in related areas concerned with the design of computer systems supporting working activity. Many within HCI and CSCW have taken up Suchman's concern with work settings and the detail of everyday working practices; as they have taken on board Suchman's arguments, they also have taken on, perhaps unwittingly, an ethnomethodological influence.

3.2. Participatory Design

One particular group that has been particularly influenced by Suchman's (1987) *Plans and Situated Actions* deserves especial mention. For a long time, there has been, within HCI, a strong and vocal group that has argued consistently that the requirements for technology should be developed directly around the work situation of the technology's users. The Participatory Design movement, in particular, has made considerable strides in developing methods and perspectives on interactive systems design from this position, for both practical reasons (concerning the efficient and fluid accomplishment of work and supporting the acceptance of technology) and political ones (emphasizing the importance of the worker's voice in issues of workplace management and development).

Because ethnomethodology generally is concerned with the "detailed and observable practices which make up the incarnate production of ordinary so-

cial facts" (Lynch, Livingstone, & Garfinkel, 1983, p. 206), its investigations of particular work domains contain rich descriptions of work practice. They begin with what is involved in the everyday accomplishment of work, not in abstract models of the activity, whether those abstract models are laid out by "professional sociology" or by management. This may suggest to those already concerned with the relationship between work and design that they have an "analytic ally" in ethnomethodology, that it can provide a methodological warrant for a primary concern with the details of work practice in the design of new technologies. So ethnomethodological perspectives, and ethnographic field techniques, found a receptive audience in this community, who found in it a resource for methodological sustenance or even empirical descriptions of work.

3.3. Ethnography

The rich descriptions of work on which ethnomethodological studies often are based have similarly played a role in making ethnomethodology (or ethnomethodologists) appealing to those seeking to ground the design of interactive technologies on studies of the performance of work. Ethnomethodology is here adopted as part of a general concern with the use of field investigations in design (Plowman, Rogers, & Ramage, 1995).

Of course, as we described already, ethnomethodology is scarcely the first form of sociological investigation to make use of ethnographic or other field techniques in working settings. The Chicago School of sociology, which emerged in the 1920s, used ethnography in turning an anthropological eye not to the tribes of the South Pacific but to the life of American cities, and so ethnography has become a technique that has been applied widely to the studies of work settings from the perspective of technological design and evaluation.

So, drawing on the sorts of confusions that we attempted to resolve earlier, it is possible that ethnomethodology comes to HCI and CSCW research, in part, in the guise of experienced fieldworkers; thus, to an extent, ethnomethodology may be basking in the sun of ethnography.

3.4. HCI in Transition

Finally, Grudin (1990) argued that HCI has passed through a number of stages in its development to reach its current position. In his characterization, it is currently moving from the fourth stage, which is focused on a dialogue with the user, to a fifth stage. This fifth stage that we are approaching is one focused not around the individual, but around the work setting.

From this perspective, we can, again, see that ethnomethodology's concern with the organizational situatedness of work might be appealing to many con-

cerned with HCI design. Ethnomethodology may be, for some, a port in the storm of transition. With its focus on the setting and situation of working activity, it may be seen as offering candidate solutions to the problems of incorporating into HCI design understandings of work setting as well as work practice.

4. ETHNOMETHODOLOGICAL STUDIES OF TECHNOLOGICAL WORK

As we stated, a variety of studies have applied ethnomethodological understandings to the lived experience of work with technology, and have been used, in turn, to support the development of new technologies and new approaches to computational support for work. As might be expected, a focus on the variety of ways in which the sequential organization of working activities are organized and the detailed practices by which they manage their work has been a common concentration in these studies.

Our goal here is not to reproduce the detailed findings of these studies. Instead, we are concerned with the relationship between the disciplines of ethnomethodology and computer systems design that these investigations embody. In particular, our focus in this section is on how ethnomethodological understandings make their way into novel system designs—on how system design “learns” from ethnomethodology.

4.1. Learning From Ethnomethodologists

To date, the most numerous examples of ethnomethodologically informed system design have been conducted by bringing together ethnomethodologists and computer scientists in multidisciplinary design teams. The investigation of Air Traffic Control by a group from Lancaster University (Bentley et al., 1992; Hughes et al., 1993) exemplifies this approach and shows how valuable it can be. In this approach, a disciplinary division of labor typically emerges. Ethnomethodologists are sent into the field and return brimming with observations and an analytic frame within which to interpret them. These observations become requirements for the system design process, more or less formally. The ethnomethodologists typically also will be involved in the ongoing evaluation of design alternatives, acting as proxy for the end-users, or, perhaps more accurately, as proxy for the work setting itself. Otherwise, they hand their requirements to their computer science colleagues, who then build the system on this basis. The fact that these requirements are derived from particularly ethnomethodological studies and understandings is all but invisible to the system developer.

This approach, and the division of labor it sets up, has become so paradigmatic that it has even driven others to frame their work as if it were being con-

ducted in this way, even when the setting is quite different. Plowman et al. (1995) observed that, to be published in the CSCW literature, it is almost required of field study reports that they conclude with a section on Implications for Design and a set of bulleted points framing observations as requirements for design, whether or not the study is actually conducted as part of an explicit design effort.⁸

Because our concern in this article is with the disciplinary relationship between ethnomethodology and computer science in this work, it is reasonable to ask, Where is the ethnomethodology in this process? In what way has ethnomethodology come to be integrated into the system design process? And what aspects of design have changed as a result?

Perhaps surprisingly, given the role of ethnomethodology in accounts of this form of system design (Sommerville, Rodden, Sawyer, & Bentley, 1992), ethnomethodology (*per se*) does not seem to have entered the process at all, save in one way—that part of the process is conducted by ethnomethodologists. The locus of ethnomethodology in this approach, then, is the ethnomethodologists's head; it consists in how part of the process (requirements capture) is conducted. Otherwise, all remains as it was; the process has not changed, and neither have the artifacts it produces.

4.2. Learning From Ethnomethodological Accounts

Less numerous, but still significant, are studies in which there is a greater disconnection between the ethnomethodological work and the system design. In these cases, the implications and requirements for design are not drawn directly from the ethnomethodologists' interpretation of their fieldwork but from the ethnomethodological accounts of such studies.

Some of our own earlier work at the Rank Xerox Research Center⁹ illustrates this approach. Bowers et al. (1995) reported on an investigation of the use of workflow technologies in the production printing industry. They described the use of a particular technology at various sites of Establishment Printers, one of the United Kingdom's largest production printing organizations. In particular, they detailed first the ways in which the model of the printing process embodied by the technology and its relationship to the management of the work systematically undermines the practices by which

8. It seems ironic that, in disciplines that have (quite rightly) rejected the right of system designers to set themselves in the place of psychologists and sociologists, the right of psychologists and sociologists to set themselves in the place of system designers seems assured.

9. Formerly Rank Xerox EuroPARC and now the Xerox Research Centre Europe.

the print workers manage the flow of work through the print shop, and second, the variety of ways in which the print workers undermine the technology to get the work done. The disparity between work process (represented explicitly within workflow technologies) and work practice (the detailed ways in which the process is actually performed) is a common focus of ethnomethodological studies of work and is highlighted by their observations.

Subsequently, the Freeflow project (Dourish, Holmes, MacLean, Marquardsen, & Zbyslaw, 1996) focused on the design of workflow technologies that would be more sensitive to the variety and fluidity of work practice. Freeflow separated the sequential logical order of working tasks from the sequential temporal order and allowed greater flexibility not only in the specification of process descriptions but also in their enactment. This was a systems design project, exploring new conceptual and architectural approaches to the design of workflow systems, but at its heart was an attempt to resolve precisely the sorts of problems that Bowers and colleagues had uncovered, both abstractly and in particular.

This project exemplifies this second use of ethnomethodological investigation as a basis for the design of new technologies. The locus of ethnomethodology is the account of work. In many ways, this can be seen as a more satisfactory way to proceed, at least from the point of view of disciplinary connection; after all, ethnomethodology (as opposed to other analytic perspectives on social interaction) becomes a part of what is communicated. On the other hand, the "disconnection" between work site and design, which is implied by this approach, also can be problematic, because it undermines, to an extent, precisely the sort of motivations that have led us, in HCI, to a deeper concern with the users of technology. If the user remains a "scenic feature of the design space" (Sharrock & Anderson, 1994), then it scarcely matters whether that feature is painted with an ethnomethodological brush.

5. ETHNOMETHODOLOGY FOR CRITIQUE AND DESIGN

Taken together, ethnomethodologically affiliated studies have produced a strong critique of the design of technology at work. They have displayed the fact that technology, at best, often fails to support the work it is designed for and, at worst, does not allow people actually to engage in their work, because the technology is not aligned to the practices through which they organize their actions, interactions, and work. Heath, Jirotko, Luff, and Hindmarsh (1994) summarized this conclusion for CSCW:

Despite impressive technological developments in CSCW, it is widely recognised that there are relatively few examples of successful applications in real

world settings. ... It is suggested that the lack of success of CSCW systems derives not so much from their technological limitations, but more from their insensitivity to the organisation of work and communication in real work environments. (p. 147)¹⁰

Suchman's (1994) discussion of the technical and theoretical basis of Winograd and Flores's "The Coordinator" provides a salutary case in point. Suchman's telling analysis of the use of speech act theory in this system, based in part on CA and the work of Sacks, emphasizes the ways in which the stipulative organization imposed by the system undermines the interactionally contingent aspects of language use; yet at the same time (and as is plaintively acknowledged in some commentaries on her article appearing in the same issue), it would seem almost to leave the practice of technological design with nowhere left to go.

5.1. Two Paradoxes of Ethnomethodologically Informed Design

Ethnomethodological analyses have been used in a range of circumstances to critique technological design in particular working settings and situations. Ethnomethodology, in attending in particular to the details of everyday action and work practice, has been able to expose an unfortunate paradox in the design of technologies for collaborative activity (or socially constructed action). This is the paradox of system design—that the introduction of technology designed to support "large-scale" activities, although fundamentally transforming the "small-scale" detail of action, can systematically undermine exactly the detailed features of working practice through which the "large-scale" activity is, in fact, accomplished. Fundamentally, it points to the interdependence of minute practice and grand accomplishment.

However, in so doing, ethnomethodology finds itself caught in a second paradox—the paradox of technomethodology.¹¹ Given the concern with the particular, with detail, and with the moment-by-moment organization of action, how can ethnomethodology be applied to the design of new technologies? Certainly, ethnomethodologists have urged that designers take into account the methods and practices through which social action, interaction, and categories of work are organized, but in the face of the unavoidably

10. More expansively, Cooper and Bowers's (1995) discussion of the "disciplinary rhetoric of HCI" pointed to the way in which construals of "user" in HCI embody an explicit move away from technology and from technological determinism, and so a focus on critique should not surprise us.

11. Taken together, these two paradoxes constitute what Grudin and Grinter (1995) referred to as "the ethnographer's dilemma."

transformational nature of technology and system design in working settings, it would seem that ethnomethodology becomes relatively powerless. Its tradition is in analyzing practice, rather than “inventing the future.”

5.2. Critique and Design

Whatever the historical context and the factors that shaped the emergence of particular ideas at particular times, we take it as fundamental that the crucial (or prosaic) reason for ethnomethodology's newfound place in the design disciplines is that there are a number of ethnomethodologists who are interested in design. The elementary components of everyday life that proved so interesting to Garfinkel, Sacks, and colleagues in the early days of ethnomethodological inquiry are, in the 1990s, increasingly dominated by (or at least suffused with) technology. For many people, everyday interactions increasingly include interactions with computer technology in one form or another. Our goal, then, is not simply to look at how ethnomethodology can be used to critique technologies, crucial though that is, or even to apply ethnomethodological understandings to better understand the conditions in which technology comes to be developed (as has been one focus, for instance, in the Participatory Design movement). Rather, alongside those investigations, we have been engaged in a different one: to understand how ethnomethodological understandings of human social action and interaction can be used, directly, in designing interactive technologies. Our focus is on design, not on critique; but it is also, critically, on the artifacts we design and the conceptual and technical apparatus by which design activity is performed, rather than on the design of specific systems.

6. TECHNOMETHODOLOGY: DRAWING FOUNDATIONAL RELATIONSHIPS

In contrast with the approaches to ethnomethodologically informed design discussed previously, our approach can be stated quite simply. We consider the relationship between ethnomethodology and system design in a design context to be more than a practical matter. For us, it is a matter of analytic orientation rather than project management. We take the ethnomethodological perspective on human social action to have detailed and deep consequences for what Suchman (1987) called the “problem of human-machine communication,” and we therefore see it as concerning the foundational concepts of system design, not simply the process by which system design proceeds. The interaction between our disciplines takes place in the interactions between the foundational elements of each discipline, not in their application to particular problems (although, of course, it is in the application that the interaction becomes valuable

and visible). Our goal is to draw foundational relationships from which to proceed together.

We refer to this approach as "technomethodology." Although this term is rather whimsical, its force is in how it emphasizes that this is not simply technology design, and neither is it simply ethnomethodology. Rather, it is something new, equally radical in its consequences for its "parent" disciplines, but different from each.

What do we mean by "drawing foundational relationships"? The answer is best given in terms of examples, and much of the rest of this article will provide an extended example of the technomethodological approach, drawing on a relationship between the notion of "abstraction" in system design and the notion of "accountability" in ethnomethodology. These are examples of the foundational elements on which we draw, elements that are conceptually central to the disciplines. Abstraction is the very stuff of system design; accountability is one of the primary elements of ethnomethodological reasoning. Similarly, concepts of number, identity, grouping, membership, formalization, and stability are grist to the technomethodological mill.

Fundamentally, then, our concern is not with the findings of particular ethnomethodological studies of working settings, but with the analytical frame within which those studies are conducted. Similarly, we are not immediately concerned with the design of this system or that system, but with the design of systems. System designers learn from ethnomethodology, not from ethnomethodologists or their observations; ethnomethodologists learn from computer science, not from computer scientists or their applications.

7. ABSTRACTION, ACCOUNTS, AND ACCOUNTABILITY

One example of the technomethodological approach currently under investigation is a user interaction model we call "accounts." This research is founded in a reevaluation of the role of abstraction in designing interactive systems, and that reevaluation is grounded in ethnomethodology's concern with the accountability of practical social action as discussed previously.

In this section, we begin by outlining the role of abstraction in traditional system design, and we subsequently introduce a novel architectural approach (called Open Implementation) that aims to address some practical, technical problems that arise in its use. These problems, traditionally tackled at the level of software infrastructure development, also arise in interactive systems. We employ the ethnomethodological perspective not only to consider the impact of these problems as interactional issues but also to consider how the technical solution provides an opportunity to fuse ethnomethodological understandings with an interaction design perspective.

7.1. Abstraction in System and Interface Design

Abstraction is the most fundamental tool of system design. Abstraction allows systems to be considered at different levels of detail, to be broken down into individual components, and to be reassembled again. The act of systems design is the creation and manipulation of abstractions. User interface behaviors (e.g., file copying, printing, and selection) are abstractions over the behavior of the programs that they control; the programs are sets of abstractions (e.g., procedures, arrays, and loops) that programmers manipulate to control the computer; and even our views of computers are couched in terms of abstractions (e.g., instruction sets, memory architectures, and bus interfaces) over raw transistors and electrical pulses. Even at that level, we cannot escape the abstraction of binary signals, imposed over continuous voltages.

Abstractions help us manage complexity by allowing us to hide it selectively. In systems design, abstractions typically function as “black boxes.” They are defined by the nature of their interactions with the outside world (human users or other pieces of code—the “clients” of the abstraction), which typically are defined in terms of the available functionality, procedure call conventions, and return values—what we typically refer to as the “interfaces” to the abstraction. The system’s internal mechanisms, which describe and control how it goes about doing the work it does, are intentionally not available to inspection. By hiding mechanism in this way, the two main uses of abstraction in system design can be achieved. First, systems can be built in terms of complex components with simple interfaces, rather than in terms of basic, raw mechanism (allowing us to build systems like spreadsheets out of mathematical packages and interface packages). Second, systems with the same interface can be regarded as equivalent (as is the case with programming language compilers or microprocessors).

In user interface design, the same models of abstraction show through. Human users interact with abstract interfaces to the system’s functionality (such as a print dialogue, or a direct manipulation view of a file system), which provide simple, consistent interaction by hiding the complex realities of the system mechanisms (creating a Postscript file and sending it to the printer, or copying files from a local disk to a server across a long-distance network connection). The goals of interface abstraction are similar to those of systems abstraction—reduction of complexity, modularity, consistency—and arise out of the use of similar techniques. Of course, because interactive systems are computational, the use of abstraction techniques in user interface design derives from their use in computational systems generally.

One use of abstraction in user interface design is the support for metaphoric interaction. Metaphors are brought into being by drawing equivalences between two abstractions, and they function precisely through the hiding of

mechanism (because the metaphors generally do not apply in terms of the mechanism). Breakdowns tend to occur where the mapping no longer holds, or where the details that the abstraction hides become suddenly relevant (such as when the network suddenly makes its presence felt by introducing delays into file copying, or when inserting a floppy disk suddenly makes it clear that the Mac trash can is not a concrete entity but is actually an abstraction for different sets of deleted files on different disks). However, it is significant that everyday interactions with the physical and social world notably do not display the same sort of "information hiding" characteristic (and hence tend not to exhibit these sorts of failure). Real-world machines produce noises and respond to physical interference, and their physical embodiment allows us to perceive their operation and even sometimes become involved in it. The real world is always available to be pushed and prodded to explore how it works, and human actors allow us to query their actions and motivations.¹² In other words, we organize our actions not simply around abstractions of possible action, but around the detail of the production of action and behavior in particular circumstances.

7.2. Abstraction and Glossing Practices in Everyday Interaction

This aspect of human action has been a central element of the ethnomethodological study program, and ethnomethodology's analysis offers lessons for the reconsideration of abstraction in interactive system design. Recall that, for ethnomethodologists, the critical property underpinning rational social action is accountability. Accountability is the property of action being organized so as to be "observable and reportable." For ethnomethodology, the key element of rational social action is that it is organized so that it can be rationalized. Garfinkel and Sacks (1970) used the notation "doing [being a researcher]," say, to describe not only the performance of the research, but (the work of) doing it so that it is (organized to be) recognizable to others as research.

Critically, this metawork—the organization of activity in this way—is not done externally to the activity itself but is, rather, a phenomenon of the activity. Activity itself is made observable and reportable, rather than being pointed to, observed and reported on. For the purposes of HCI, then, the critical observation is this. What computational abstractions share with the abstractions of natural, everyday interaction is that they are organized to reveal certain things (and hide others) for certain purposes. What they do not share with the abstractions of everyday activities is the observable–reportable na-

12. This is what makes those cases where we cannot ask questions (interactions with "faceless bureaucracies") so frustrating.

ture of everyday action that is at the heart of ethnomethodological investigation. Garfinkel and Sacks (1970) referred to the mechanisms at work here as *glossing practices*—“methods for producing observable–reportable understandings, ... a multitude of ways for exhibiting-*in*-speaking and exhibiting-*for*-the-telling that and how speaking is understood” (pp. 333–334).

In other words, the key property of human action is the way in which it is made observable and reportable in the course of its own unfolding. It is organized accountably. As someone speaks, this is how “he provides the very materials for making out what he says” (Garfinkel & Sacks, 1970, p. 344). Computational abstractions, being static, atomic, and unexaminable, provide no such means. Abstract computational behavior is not accountable; for this reason, the forms of contingent, improvised (situated) action become problematized.

Are we simply saying, then, that ethnomethodology suggests that it is a good idea to design computer systems so that people can understand them? That would hardly be news. Making systems understandable, less inscrutable, and more open to examination has, after all, been the primary focus of HCI for all these years. But, of course, we are saying more than this. What ethnomethodology tells us is that the production of an account of action is an indexical (or situated) phenomenon. In other words, a user will encounter a system in myriad settings and circumstances and will attempt to find the system’s behavior rational and sensible with respect to whatever those infinitely variable circumstances might be, day to day and moment to moment. What this implies, then, is that the creation of an account for a system’s behavior is not a “one-off” business. It cannot be handled once and for all during a design phase conducted in the isolation of a software development organization in Silicon Valley. The creation of the account happens, instead, in every circumstance in which the system is used, because the account and the circumstance of the use are intimately correlated. In technical terms, an account is a run-time phenomenon, not a design-time one.

7.3. Toward Observable–Reportable Abstractions: Open Implementation

What this leads us toward, then, is a way of opening up the abstractions that the system offers so that they can convey aspects of the mechanisms that lie behind them. By revealing more of what lies behind them, these more “translucent” interfaces would provide cues as to not only what the system was doing but why it was being done and what was likely to be done next, uniquely for the immediate circumstances. A view onto the mechanism can provide a context to make system activity rationalizable and to do so within the circumstances of its activity rather than outside of it (as, say, manuals might do). Critically, this is not a call for more complex, more mechanical interfaces;

those would be ones that do away with the abstractions, and we need to retain the abstractions to retain clarity, consistency, and ease of use. However, like social interactions in everyday life, we would like our interactions with computational systems to be organized in terms of abstractions that are supported by their own unfolding, rather than opaque (and brittle) black boxes.


This is not to say, of course, that opening up mechanism leads directly to a user's complete understanding of what is going on. There are clearly many details to be accomplished in implementation that are essentially irrelevant to the work that the abstraction does. Precisely how these details of mechanism work their way into the interface is a design issue that we address more directly shortly. Our point here, simply, is that to manage the relationship between the user's work and the system's action more effectively, we need to provide users with more information about how the system goes about performing the activities that have been requested and that the place to look for this information is within the implementation, below the abstraction barrier.

A recent line of development in software architecture, for quite different reasons, has been moving in just this direction. Open Implementation (OI; Kiczales, 1996) is an approach to system design that recognizes that the implementation features that abstractions hide are often ones that embody design decisions critical to the effective use of the abstraction. The very notion of abstraction that supports modularity and reuse in system design also may make modules harder to reuse. Kiczales' (1996) concern is with software abstractions offered to programmers and to other system components, but precisely the same arguments apply to interactive systems, and the techniques that OI develops are also applicable in this domain as one possible route to providing accounts of system action.

One key principle that underlies much OI practice is computational reflection (B. Smith, 1982). Reflective systems are computational entities that contain a representation of aspects of their own structure and behavior. Critically, this representation is causally connected to the behavior that it describes. The result of this causal connection is that not only will changes in the system's behavior be reflected by equivalent changes in the representational model, but the model can be changed to change the system. This principle originally was applied in the design of programming languages, providing languages with models of their own execution that could be used to introduce new programming features in an implementation-independent way. The OI approach has applied it more widely to the problems of abstraction, using a reflective "metalevel" to offer a model of the internal mechanism lying inside an abstraction, as a means to observe and control how that abstraction will be realized when the system is used (see Figure 1).

The full detail of the OI approach and, in particular, the practicalities of implementing reflective systems are beyond the scope of this article; the reader is referred to more technical descriptions such as that of Kiczales, des Rivieres,

Figure 1. (a) Client's interaction with traditional black-box abstractions through standard abstraction barriers. (b) Open implementations also reveal inherent structure.



and Bobrow (1991) for a fuller account of the techniques by which OIs are realized. Here, however, we sketch relevant aspects of OI's reformulation of computational abstraction.

The goal of OI is to provide flexibility and the opportunity for reconfiguration of the mechanism that lies behind a traditional computational abstraction. Drawing on examples of current practice in a wide range of domains, OI argues that details of implementation strategy—details that are explicitly hidden by traditional abstraction techniques—will often be critical to the way in which clients will make use of the abstraction. Traditional abstractions, embodied in programming languages, Application Programming Interfaces, and user interfaces, talk only in terms of input and output, request and result, but never in terms of how the input will be mapped to the output or how the result will be generated when the request is made. If a programmer knows that creating a new window in a window system is a fast operation with low memory overhead, then she can program her application to use windows whenever a new screen object is required. On the other hand, if she knows that there are significant performance and memory overheads involved in creating windows, she might choose instead to allocate just one large window for her application and handle the allocation of space inside it herself. The traditional abstraction model supporting most software provides no opportunity for gaining this sort of information because it is locked within the "black box".¹³

13. That is, if it is to be found anywhere at all. Some concerns may be manifest only when the system is running and have no place in the system's own implementation at all. For example, most network software has some way of adapting to the current traffic level on a network, but most are unlikely to have a direct measure of what the traffic level actually is.

The OI strategy is to provide two interfaces to the implementation, at different levels. The first is the traditional interface, by which clients can make use of the abstraction (e.g., create, draw in, and destroy windows). This is called the *base-level interface*. In addition, it provides a second interface, called the *metalevel interface*. The abstraction offered at this interface is a rationalized model of the inherent structure of the implementation, and the controls offered at the metalevel interface can be used to control aspects of the implementation of the base-level interface; that is, it provides the programmer with an abstract view into the mechanism, through which aspects of that mechanism can be controlled (see Figure 1).

This approach to computational architecture has been fruitfully applied in a number of areas of system design, most notably in the design of programming languages (e.g., the Common Lisp Object System; Bobrow et al., 1988) and operating systems (e.g., Apertos; Yokote, 1992). Briefly, its value is that, in opening up to scrutiny how traditional abstractions are to be realized, it allows programmers to make critical distinctions between what they want to do with system abstractions and how they want the system to go about providing its functionality. Furthermore, the approach does this in a way that makes clear the relationship between what and how; they are not divorced from each other.

For our purposes, then, the critical feature of the OI approach, and the feature on which we rely in developing its use as in interactive systems, is the way in which it allows us to forge and articulate the relationship between what is done (the implementation behavior), and what is done by what is done (the achievement of application ends).

7.4. Accounts and Accountability

We currently are developing an approach to the design of interactive systems in terms of OI. One aspect is to adopt the OI model in the same way in which it has been applied to other domains: as an engineering technique that can be used to provide considerably greater flexibility than would be available otherwise. For instance, Prospero (Dourish, 1996) is a CSCW tool kit that employs OI in this way, offering programmers the opportunity to become involved in the implementation of infrastructure mechanisms that support their applications.

We are concerned here with a second approach. This is a more fundamental one and aims to address the disparity between traditional process-driven models of interface design and the more improvisational model revealed by sociological and anthropological investigations such as Suchman's. This approach is based on a rereading of OIs' reflective self-representations as accounts that systems offer of their own activity (Dourish, 1997).

As should be clear from what has gone before, the term *account* is chosen particularly to emphasize a metaphorical frame drawn from the ethnomethodological perspective on the organization of action. So what is important about this approach is not the account itself (the explanation of the system's behavior) but rather accountability in the way this explanation arises. In particular, the account arises reflexively in the course of action, rather than as a commentary on it, and concerns the way in which that action is organized so that it can be made rational in particular circumstances. These features, which arise directly from the OI approach, allow us to use these accounts as a means for users to rationalize the activity of the system and therefore to organize their behavior around it, as interaction proceeds, for their own practical purposes.

8. EXAMPLE: ACCOUNTING FOR FILE COPYING

To make these ideas more concrete, we present an example of working with accounts in this section. The role of the example is to illustrate not only the role of accounts in interactive systems, but how accounts operate as an example of applying the technomethodological approach to the relationship between abstraction and accountability.

8.1. The File-Copying Scenario

The file-copying scenario is familiar to users of graphical desktop environments, in which a single "folder" or "directory" abstraction is used for all containers of files. Folders can be specified as targets for copying operations simply by dragging the icons for files to be copied and dropping them on the destination folder. In many such environments, initiating a copy operation of this sort will cause a status indicator such as a "percentage-done bar" (or "thermometer") to appear, indicating how much of the operation has been completed. So, graphical interaction (drag and drop) initiates the copy operation, and the percentage-done bar reports progress. As more of the files are copied, more of the bar is filled in, until eventually all the files are copied and the bar is completely filled.

However, consider an alternative scenario that is possible. Imagine that the folder to which the files are copied is not actually a folder on your local hard disk but rather is a folder on a remote volume accessible over a network. Let us consider a case in which someone drags a large set of files onto the folder. A percentage-done bar appears and starts to fill, but when it reaches 40% complete, the copy operation fails. There are many reasons why it might have failed: The remote volume may have become full, or the server may have become unavailable, or the network may have failed (or perhaps it never was connected in the first place).

In this case, what resources are available to the user to understand what has happened and to understand what options are now available? What does it mean that the percentage-done bar filled to 40%? Were 40% of the files written onto the remote volume? Were 40% of the files read from the local disk? Did all of the files get 40% of the way to their destination?

From a technomethodological perspective, consider why these questions are important and how they arise. Ethnomethodology talks about the way in which people find, within the circumstances of action, the means to find that action rational. Furthermore, in terms of Suchman's work, these resources provide the means to articulate abstract plans of action and to organize the specifics of action. However, the abstraction that has been offered by the system—the folder—hides the details on which such understandings could be based. The differences between local and remote folders, the difference in the operation of local and remote copy operations, and the consequences of these differences are hidden from view.

Furthermore, it is not sufficient simply to offer two different kinds of folders, providing a distinction between local and remote. Why is this not an adequate solution? Recall our earlier observation that an account is a run-time phenomenon. Actions and accounts are situated within the specific circumstances of their production, not within abstract characterizations of them. In other words, what is important here is not the differences between two abstract types of copying (local copying and remote copying), but the specifics of this or that copying operation. There are far too many different features of the occasion (including distance, available network bandwidth, other people's activities, the types of files involved, and even the type of network infrastructure) for designers or users to be able to distinguish among them in the abstract model that the system offers. Even if we were to provide a hundred different sorts of folder copy destinations, they would remain disconnected from the actual process of copying that is taking place. Accounts, both in the strict ethnomethodological sense and in the metaphorical technical sense that we are developing here, must arise in response to and be organized around the specific circumstances of their production, which are the specific circumstances within which the action takes place.

8.2. An Account of File Copying

Instead of trying to provide different abstractions for all the different circumstances in which copying might take place, the accounts model provides a mechanism for dynamically relating manipulation of the abstraction to the detail of what is actually happening. The account provides a means for users to see what copying means in this case, by providing a view into the mechanism by which copying is carried out.

What is needed in the first instance, then, is a model of the copying mechanism—an account of file copying, in terms of which a specific account of a specific copy can be formulated. In seeking such an account, we look for the structural properties of the system's behavior with reference to which of those elements that concern us can be made manifest in the interface. The "account" here (i.e., the interface's representation of the system's behavior) is not an explanation as such, but a backdrop against which the system's behavior can be played out and made understandable.


As an example, we can characterize the structure of file copying as shown in Figure 2. Between the file source and destination are arrayed a number of staging posts (data buckets). File data flows from the start point to the endpoint by moving from one bucket to another along a data path. As data flow from one bucket to the next, the buckets are related to each other by flow strategies, by which the movement of data from one to the next is regulated. The structure of a flow path and the strategies used will vary in different circumstances; the number of buckets and the variety of strategies by which data flow is controlled characterize the different ways of doing file copying. So, in different situations, data buckets may be used to describe file system caches, network interfaces, and networks themselves. The flow of data through these and the activation of the flow strategies provide a framework for the relationship between the action in which the system engages and the reading and writing of data files.

In addition, this structure also provides the means to answer the sorts of questions that were raised earlier when the file-copy failure was observed for network copying. There were two particular problems when those questions arose. First, the system had no means to draw distinctions between the different circumstances in which file copying might be done. The data bucket and flow strategy mechanism provides just this opportunity. Second, the system provided no terms in which the role of the percentage-done bar played, making it impossible to interpret its behavior. However, with this structure in place, we can describe "where" in the flow the percentage-done bar is connected and thus allow the user to make some sense of what it is actually reporting.

8.3. Accounts and Ordinary Operation

Although our scenario described a failure in network copying (and so a requirement to provide more information to allow the user to understand what had gone wrong), it is important to note that, by being offered within the action rather than from outside it, these sorts of interface accounts provide not just for recovery from failure but also for more detailed ongoing monitoring of action. It is from this provision of information in the course of activity that the system supports the sorts of improvised activity that Suchman brought to the analysis of behavior at the interface. By enriching the system's detailing of the circum-

Figure 2. A structural model of the file-copying example in terms of data buckets and the connections between them. Connections between elements of this model are the points at which strategies and policies can be identified.



stances in which it is acting, we similarly enrich the potential resources for the user's moment-by-moment decision making.

In the case of file copying, an explanatory system organized around failure would be useless to make decisions like "Why is this taking so long?" or "Will this finish before my ride home arrives?"—the sort of questions that potentially lie at the heart of decisions to stop the copy, to do it in another way, to copy a subset of the files, and so on. In other words, an explicit failure model sets limits not only on the sorts of questions that might be asked, but also, in organizing them around specific breakdowns, on the reasons that those questions might arise. Our goal is to avoid this problem by providing information without making a prior commitment to the reason that information might be useful. Such information is useful in cases in which there is no technical failure, or even no failure at all.

The account, then, is not simply a new form of error-reporting system, but rather becomes part and parcel of ordinary interaction with a computer system. Again, this draws on the ethnomethodological perspective on interaction, where accountability arises not out of specific requests for information but, first, as part and parcel of everyday activity and, second, as a crucial feature allowing concerted action to arise.

An account, too, is partial. It reveals certain features and hides others; in fact, what makes this an account of file copying is that it talks in terms of files and copies and says nothing about (say) DMA, disk blocks, and SCSI adapters. On the other hand, it is clear that one account leads to another; even in our simple model in Figure 2, we are led to consider the "name mapping strategy" (i.e., the fact that at the source end, a name points to a file that already exists, and at the destination end, it points to one that will only exist when the operation has finished). So we live within a world that is endlessly accountable

in different terms and to different ends. What is critical here is the way in which the reflective model provides a basis for enriching computational action with structural accounts of what that action is.

8.4. Accounts and Mental Models

On brief inspection, there might appear to be considerable overlap between the accounts-based notion of interaction offered here and the body of work on mental models in HCI (see, e.g., Gentner & Stevens, 1983, for an introduction and overview of this work). It is useful to explore why this is not the case.

The fundamental distinction is quite straightforward. A *mental model*¹⁴ is a model of the operation of a system, formed by a user of that system, that users employ in the planning of their activity. An *account*, on the other hand, is a model of the system's activity offered by the system to account for and cast light on its own action. Although this distinction is easy to state, and the fundamental difference in the "location" of the model is clear, it hides a range of other issues. Two concern us here.

The first issue is that, as can be seen in the phrasing of the terms previously discussed, the traditional notion of mental model, as derived from a cognitivist perspective, originates from a very different position on human action than that which we are considering here. This is not the place for a debate on the relative merits of phenomenological sociology and cognitive science. However, it is important to emphasize the way in which the accounts model arises from a model of situated, resource-driven activity rather than one constituted in terms of abstract models of behavior. It is this same perspective that leads us to frame the account in terms of what the system does rather than what the user might have meant.

The second issue is that, as part of the artifact, the account is itself designed along with the rest of the system. It is part of the interface (albeit at a metalevel, rather than concretely instantiated in buttons and menus). The questions that this observation raises—of how the artifact is designed, of the principles for its appearance in the user interface, and of the means by which its structure is determined—are subject to ongoing investigation (e.g., Dourish & Curbow, 1997). Our goal, however, is not to develop a design method by which the "right" account can be found and embodied in the system, but rather to investigate this "two-level" approach to interface design as a way of making

14. The term *mental model* has been used by different authors at different times to refer to different things. We take this meaning, a conceptual model of a system possessed by the user of that system, to be the primary meaning of the term, and the one most relevant here.

ethnomethodological perspectives on accountable action "real" in the abstractions of interactive systems.

8.5. Accounts as an Example of Technomethodology

The notion of accounts, as described here, is an example of the technomethodological approach. It is an approach to the design of interactive systems, based on a reconsideration of a fundamental concern in systems design, that of software abstraction. At the same time, that reconsideration is based on the foundational analytic principles of ethnomethodology. In doing this, we achieve our twin goals of providing ethnomethodological insight for the design of interactive systems and computational leverage for the practice of ethnomethodologically inspired design.

We principally have been developing our ideas using this model of accountable interactive systems as a foil, and we have dwelt on it here because it was necessary to discuss an example of our approach to "foundational relationships" in some detail. It is important to remember, however, that this is only a single example.

We currently are investigating further opportunities for foundational relationships between ethnomethodological ideas and the basic principles of computational design. For instance, much of computational practice is based on a notion of type (or class), as well as complementary notions of grouping (sets, bags, lists, arrays, etc.). These ideas, and the notions of identity on which they are based, are sufficiently fundamental and pervasive that they are typically manifest throughout a system, right up to the user interface. Social categories, however, have some quite different properties, being typically more fluid and relative to individuals and tasks. We are interested in the opportunities for an interactive system to exploit a form of identity and grouping that has a stronger grounding in the practical exercise of membership conditions in everyday action, as revealed in ethnomethodological investigations (e.g., Sacks, 1972). Recent explorations of richer models of identity and membership, such as Chambers's "predicate classes" (Chambers, 1993; Ernst, Kaplan, & Chambers, 1998) or "subject-oriented programming" (Harrison & Ossher, 1994; R. Smith & Ungar, 1996), may serve as the basis of new approach.

In other words, there are a range of opportunities for the sorts of foundational relationships that drove our investigation of accountability and for which we have argued here. Although accounts is one example, our goal here is to motivate the general approach. To that end, before we close, we address one remaining general issue concerning the relationship we propose between ethnomethodology and interactive systems design.

9. ABSTRACTION AND PARTICULARITY

The reported phenomena are only inspectably the case. They are unavailable to the arts of designing and interpreting definitions, metaphors, models, constructions, types or ideals. They cannot be recovered by attempts, no matter how thoughtful, to specify an examinable practice by detailing a generality. (Garfinkel, 1991, p. 16)

Our approach seems, on the surface, to be in contradiction to some widely held views—perhaps even, paradoxically, the fundamental tenets of the disciplines we are trying to relate.

In CSCW, particularly, where the disciplinary constitution of the field means that these deliberations arise frequently, ethnomethodologists are characterized, somewhat unfairly, as unwilling or unable to traffic in the generalizations that system design needs. Computer scientists, in turn, sometimes are seen by ethnomethodologists as overeager to generalize findings, stripping them in the process of precisely the finely balanced features of the setting that are the very essence of the ethnomethodological accounts. To those who are familiar with these oft-repeated observations (whether printed in the pages of learned journals or grumbled around the bar), we would seem to be setting ourselves an impossible task—to exploit generalizations in ethnomethodology rather than working from particular field studies and, what's more, to exploit these generalities not in terms of particular application settings, but in terms of the conceptual underpinnings of systems design.

The resolution of this problem lies in a distinction between the ways in which our two disciplines use abstraction. For computer science, abstractions are generative—that is, they are used to generate behaviors. There is no way to put a character on a screen without manipulating the abstractions of windows, fonts, and display procedures. The results of much of what computer scientists do every day (building architectures, descriptions, and programs) are based on the application and manipulation of abstractions.

Despite appearances to the contrary, ethnomethodology also traffics in abstractions, but they are abstractions of a different sort. For ethnomethodology, abstractions are analytic, used to characterize and explain behavior (but not necessarily to take an active role in how that behavior comes about). What ethnomethodology rejects is not the notion of widely applicable abstractions and concepts, but rather the “ideal types” of traditional sociology. Ideal types, a concept first put forward and exploited by Weber, strip away distracting detail to reveal the “essential patterns” of social truth underneath. In effect, they provide the sociological theorist with a means to step outside the vagaries of everyday concrete phenomena in providing a sociological account of the world, the very phenomena that ethnomethodology embraces and revels in.

Ethnomethodology rejects abstraction only inasmuch as it rejects an approach to sociological theorizing that deals with the problem of social order in terms of ideal types and social structure.

In other words, there is a considerable difference between the abstractions of conventional sociology and those of ethnomethodological investigations, in terms of their ontological status. Whereas conventional sociology sought the elements of a theoretical mechanics of social order, ethnomethodologists sought to uncover only those regularities that they could observe directly in naturally occurring data.

The power of those regularities, however, is precisely how they occur on different occasions, and indeed many of the social "mechanisms" that ethnomethodology has described are found across many different social circumstances. For example, Garfinkel and Sacks (1970) described "cohort independence" phenomena, by which they meant social phenomena that are not tied to the scenic features of their production. Thus, for example, the model for turn-taking in conversation (Sacks, Schegloff, & Jefferson, 1974) is, in some crucial respects, cohort-independent in that it operates across local circumstances such as gender, ethnicity, race, occupational identity, and so forth. It is these sorts of generally operative social processes, explicated by ethnomethodology, that we turn to in our technomethodological approach. These processes are, in Sacks et al.'s terms, "context free, yet context sensitive."

The crucial feature of these abstractions is that, abstract though they are, they are uncovered and operate in the particular. This concern with the particular is a fundamental methodological tenet of ethnomethodology, dealing with naturally occurring everyday behavior rather than with generalizations of social action. They are abstract in that they are not concrete; they do not set out how it is that any given conversation will unfold. But they are abstract nonetheless.

The distinction between these two dimensions—between the abstract and the concrete, between the general and the particular—provides us with the opportunity to engage in technomethodology. In particular, the move represents an attempt to work with a set of sensibilities rather than with the details of specific activity, even though, of course, those sensibilities arise out of the ethnomethodology's very concern with the grounded and specific experience of everyday activity. In other words, technomethodology attempts to align system design not so much with the details of specific working practices as with the details of the means by which such working practices arise and are constituted.

Consider an example to illustrate this distinction. In recent years, there has been an interest in using the insights of ethnomethodology for the development of dialogical interfaces. Attempts have been made to build into com-

puter interfaces the specifics of Sacks et al.'s (1974) turn-taking model, such as the rules associated with speaker transfer. However, our argument is that the value of the turn-taking model described by Sacks et al. is in the way in which it shows how the abstractions of conversational flow are sustained rather than rote procedures by which they might be enacted (Button & Sharrock, 1995b). It is this notion of the ongoing management of conversation, rather than the specifics of any human dialogue, that provides an abstraction for design. When we fail to make the distinction, we fall foul of the paradox of technomethodology.

10. CONCLUSIONS

Although the influence of research work conducted from the perspectives of sociology and anthropology has become increasingly prominent in HCI in recent years, we typically observe sociological observation rather than sociological analysis finding its way into new research systems. Ethnomethodology, which has become particularly prominent, especially in the field of CSCW, seems to disappear as an analytic position in favor of the observations that ethnomethodologically inclined fieldworkers report from their observational studies. The goal of our work, reported here, has been to seek a new position on the relationship between computer science and ethnomethodology in the design of interactive software. This position regards the relationship between our disciplines as a foundational, analytic concern rather than simply a practical one, and so it emphasizes how it is that the ethnomethodological position on the problem of social order can inform, respecify, and reconceptualize foundational elements of system design.

Our investigation of these issues takes place on three levels. First, we are concerned with developing a basic analytic position on the foundational relationship between our two disciplines. Our goal is a position that moves away from a "service" relationship between our disciplines, in which ethnomethodologists uncover requirements for system design or computer scientists develop systems organized around specific working practices. Second, we are engaged in developing a collection of specific disciplinary relations both as illustrations of our position and as tools of a new model of design. The work presented here on accountability and abstraction is one example of a relation of this sort; ongoing work is directed toward other particular relations. Third, we are looking at how these relations can be exploited in specific design projects to chart a course from analytic reasoning to working software systems. We hope to report on these investigations in the future.

NOTES

Acknowledgments. The seeds of the work reported here grew from the long-term interest in the relationship between social and design sciences, which was a core research focus at the Rank Xerox Research Centre and the Xerox Palo Alto Research Center. We acknowledge the valuable contributions, to both our thinking and writing, made by many colleagues there and elsewhere, particularly Annette Adler, Bob Anderson, Victoria Bellotti, Beki Grinter, Austin Henderson, Yvonne Rogers, Lucy Suchman, Randy Trigg, and Terry Winograd. A conversation with Lynn Cherny prompted the first written expression of these ideas.

Authors' Present Addresses. Paul Dourish, Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304. E-mail: dourish@parc.xerox.com. Graham Button, Xerox Research Centre Europe, 61 Regent Street, Cambridge, CB2 1AB, United Kingdom. E-mail: button@xrce.xerox.com.

HCI Editorial Record. First manuscript received October 29, 1997. Revision received September 26, 1998. Accepted by Terry Winograd. Final manuscript received October 23, 1998. —*Editor*

REFERENCES

- Anderson, R. (1991). Representation and requirements: The value of ethnography in system design. *Human-Computer Interaction*, 9, 151-182.
- Anderson, R. (1996). Work, ethnography and systems design. In *Encyclopedia of microcomputers* (Vol. 20, pp. 159-183). New York: Marcel Dekker.
- Anderson, R., Button, G., & Sharrock, W. (1993). Supporting the design process within an organisational context. *Proceedings of the ECSCW'93 European Conference on Computer-Supported Cooperative Work*, 47-60. Dordrecht, The Netherlands: Kluwer.
- Bentley, R., Rodden, T., Sawyer, P., Sommerville, I., Hughes, J., Randall, D., & Shapiro, D. (1992). Ethnographically-informed systems design for air traffic control. *Proceedings of the CSCW'92 Conference on Computer-Supported Cooperative Work*, 123-129. New York: ACM.
- Bobrow, D., Demichiel, L., Gabriel, R., Keene, S., Kiczales, G., & Moon, D. (1988). *Common lisp object system specification* (X3J13 Document 88-002R). New York: American National Standards Institute.
- Bowers, J. (1994). The work to make a network work: Studying CSCW in action. *Proceedings of the CSCW'94 Conference on Computer-Supported Cooperative Work*, 287-298. New York: ACM.
- Bowers, J., Button, G., & Sharrock, W. (1995). Workflow from within and without. *Proceedings of the ECSCW'95 Fourth European Conference on Computer-Supported Cooperative Work*, 51-66. Dordrecht, The Netherlands: Kluwer.
- Bowers, J., & Pycok, J. (1994). Talking through design: Requirements and resistance in cooperative prototyping. *Proceedings of the CHI'94 Conference on Human Factors in Computing Systems*, 299-305. New York: ACM.

- Button, G., & Sharrock, W. (1995a, May). *Ethnomethodology and HCI*. Paper presented at the CHI'95 Research Symposium, Denver, CO.
- Button, G., & Sharrock, W. (1995b). Simulacrum of conversation. In P. Thomas (Ed.), *The social and interactional dimensions of human-computer interfaces* (pp. 107-125). Cambridge, England: Cambridge University Press.
- Chambers, C. (1993). Predicate classes. *Proceedings of the ECOOP'93 European Conference on Object-Oriented Programming* (pp. 268-296). Berlin: Springer-Verlag.
- Cooper, G., & Bowers, J. (1995). Representing the user: Notes on the disciplinary rhetoric of human-computer interaction. In P. Thomas (Ed.), *The social and interactional dimensions of human-computer interfaces* (pp. 48-66). Cambridge, England: Cambridge University Press.
- Dourish, P. (1996). *Open implementation and flexibility in CSCW toolkits*. Unpublished doctoral dissertation, University College, London.
- Dourish, P. (1997). Accounting for system behaviour: Representation, reflection and resourceful action. In M. Kyng & L. Mathiassen (Eds.), *Computers and design in context* (pp. 145-170). Cambridge, MA: MIT Press.
- Dourish, P., & Curbow, D. (1997). *Exploring the interface to accountable systems: From direct manipulation to direct experience*. Unpublished technical report, Apple Computer, Inc., Cupertino, CA.
- Dourish, P., Holmes, J., MacLean, A., Marqvardsen, P., & Zbyslaw, A. (1996). Freeflow: Mediating between representation and action in workflow systems. *Proceedings of the CSCW'96 Conference on Computer-Supported Cooperative Work*, 190-198. New York: ACM.
- Ernst, M., Kaplan, C., & Chambers, C. (1998). Predicate dispatch: A unified theory of dispatch. *Proceedings of the ECOOP'98 European Conference on Object-Oriented Programming*, 186-211. Berlin: Springer-Verlag.
- Frohlich, D., & Luff, P. (1990). Applying the technology of conversation to the technology for conversation. In P. Luff, N. Gilbert, & D. Frohlich (Eds.), *Computers and conversation* (pp. 187-220). New York: Academic.
- Garfinkel, H. (1967). *Studies in ethnomethodology*. Englewood Cliffs, NJ: Prentice Hall.
- Garfinkel, H. (1991). Respecification: Evidence for locally produced, naturally accountable phenomena of order, logic, reason, meaning, method, etc. in and as of the essential haecceity of immortal ordinary society (I)—An announcement of studies. In G. Button (Ed.), *Ethnomethodology and the human sciences* (pp. 10-19). London: Routledge & Kegan Paul.
- Garfinkel, H., & Sacks, H. (1970). Formal structures of practical action. In J. C. McKinney & E. A. Tiryakian (Eds.), *Theoretical sociology* (pp. 337-366). New York: Appleton-Century-Crofts.
- Gentner D., & Stevens, A. (1983). *Mental models*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.
- Grinter, R. (1997). Doing software development: Occasions for automation and formalisation. *Proceedings of the ECSCW'97 European Conference on Computer-Supported Cooperative Work*, 173-188. Dordrecht, The Netherlands: Kluwer.
- Grudin, J. (1990). The computer reaches out: The historical continuity of interface design. *Proceedings of the CHI'90 Conference on Human Factors in Computing Systems*, 261-268. New York: ACM.

- Grudin, J., & Grinter, R. (1995). Ethnography and design. *Computer-Supported Cooperative Work*, 3, 55-59.
- Harrison, W., & Ossher, H. (1994). Subject-oriented programming—A critique of pure objects. *Proceedings of the OOPSLA'94 Conference on Object-Oriented Programming Languages, Systems and Applications*, 411-428. New York: ACM.
- Heath, C., Jirotko, M., Luff, P., & Hindmarsh, J. (1993). Unpacking collaboration: The interactional organisation of trading in a city dealing room. *Proceedings of the ECSCW'93 Third European Conference on Computer Supported Cooperative Work*, 155-170. Dordrecht, The Netherlands: Kluwer.
- Heath, C., Jirotko, M., Luff, P., & Hindmarsh, J. (1994). Unpacking collaboration: The interactional organisation of trading in a city dealing room. *Computer-Supported Cooperative Work*, 3, 147-165.
- Heath, C., & Luff, P. (1991). Disembodied conduct: Communication through video in a multi-media environment. *Proceedings of the CHI'91 Conference on Human Factors in Computing Systems*, 99-103. New York: ACM.
- Heritage, J. (1984). *Garfinkel and ethnomethodology*. Cambridge, MA: Polity.
- Hughes, J., O'Brien, J., Rodden, T., Rouncefield, M., & Blythin, S. (1997) Designing with ethnography: A presentation framework for design. *Proceedings of the DIS'97 Conference on Designing Interactive Systems*, 147-158, New York: ACM.
- Hughes, J., Randall, D., & Shapiro, D. (1993). From ethnographic record to system design: Some experiences from the field. *Computer-Supported Cooperative Work*, 1, 123-141.
- Kiczales, G. (1996, January). Beyond the black box: Open implementation. *IEEE Software*, 8-11.
- Kiczales, G., des Rivieres, J., & Bobrow, D. (1991). *The art of the metaobject protocol*. Cambridge, MA: MIT Press.
- Livingstone, E. (1988). *Making sense of ethnomethodology*. London: Routledge & Kegan Paul.
- Lynch, M., Livingstone, E., & Garfinkel, H. (1983). Temporal order in laboratory work. In K. Knoll-Centina & M. Mulkay (Eds.), *Science observed* (pp. 205-238). Newbury Park, CA: Sage.
- Parsons, T. (1937). *The structure of social action*. New York: McGraw-Hill.
- Plowman, L., Rogers, Y., & Ramage, M. (1995). What are workplace studies for? *Proceedings of the ECSCW'95 European Conference on Computer-Supported Cooperative Work*, 309-324. Dordrecht, The Netherlands: Kluwer.
- Rouncefield, M., Hughes, J., Rodden, T., & Viller, S. (1994). Working with constant interruption: CSCW and the small office. *Proceedings of the CSCW'94 Conference on Computer-Supported Cooperative Work*, 275-286. New York: ACM.
- Sacks, H. (1972). An initial investigation of the usability of conversational data for doing sociology. In D. Sudnow (Ed.), *Studies in social interaction* (pp. 31-74). New York: Free Press.
- Sacks, H. (1984). On doing 'being ordinary.' In M. Atkinson & J. Heritage (Eds.), *Structures of social action: Studies in conversation analysis* (pp. 413-429). Cambridge, England: Cambridge University Press.
- Sacks, H. (1992). *Lectures on conversation*. Oxford, England: Blackwell.

- Sacks, H., Schegloff, E., & Jefferson, G. (1974). A simplest systematics for the organization of turn-taking in conversation. *Language*, 50, 696-735.
- Sellen, A., & Harper, R. (1997). Paper as an analytical resource for the design of new technologies. *Proceedings of the CHI97 Conference on Human Factors in Computing Systems*, 319-326. New York: ACM.
- Shapiro, D. (1994). The limits of ethnography: Combining social sciences for CSCW. *Proceedings of the CSCW'94 Conference on Computer-Supported Cooperative Work*, 417-428. New York: ACM.
- Sharrock, W., & Anderson, R. (1994). The user as a scenic feature of the design space. *Design Studies*, 5, 5-28.
- Smith, B. (1982). *Reflection and semantics in a procedural language* (Tech. Rep. No. MIT-TR-272). Cambridge, MA: Laboratory for Computer Science, Massachusetts Institute of Technology.
- Smith, R., & Ungar, D. (1996). A simple and unifying approach to subjective objects. *Theory and Practice of Object Systems*, 2, 161-178.
- Sommerville, I., Rodden, T., Sawyer, P., & Bentley, R. (1992). Sociologists can be surprisingly useful in interactive systems design. In A. Monk, D. Diaper, & M. Harrison (Eds.), *People and computers VII. Proceedings of HCI'92* (pp. 341-353). Cambridge, England: Cambridge University Press.
- Suchman, L. (1987). *Plans and situated actions: The problem of human-machine communication*. Cambridge, England: Cambridge University Press.
- Suchman, L. (1994). Do categories have Politics? The language/action perspective reconsidered. *Computer-Supported Cooperative Work*, 2, 177-190.
- Thomas, P. (1995). *The social and interactional dimensions of human-computer interfaces*. Cambridge, England: Cambridge University Press.
- Yokote, Y. (1992). The Apertos reflective operating system: The concept and its implementation. *Proceedings of the OOPSLA'92 Conference on Object-Oriented Programming Languages, Systems and Applications*, 414-434. New York: ACM.