

Ghosts in the Network: Distributed Troubleshooting in a Shared Working Environment

Yvonne Rogers

School of Cognitive and Computing Sciences
The University of Sussex
Brighton
BN1 9QH
UK
(+44-0273-606755) yvonner@cogs.susx.ac.uk

ABSTRACT

The implementation of networking technology in work settings offers numerous opportunities for improving the transmission of information and the sharing of resources within and between organizations. Its success in integrating distributed working activities, however, rests on how well the users of a network can coordinate their activities with respect to each other. This paper examines the communicative and interactive processes that take place when a typical breakdown occurs in a networked environment. A detailed analysis is presented which interprets the events that unfold in relation to the socio-cognitive issues of shared understanding, the transmission of knowledge and distributed problem-solving.

KEYWORDS

Distributed problem-solving, networked technologies, ethnographic analysis, breakdowns.

INTRODUCTION

The introduction of distributed computer systems within commercial organizations is becoming widespread [11]. The linking up of remote workstations, file servers and plotters together with the use of networking software means that information previously stored in separate systems can be automatically transferred between different locations. Such a capability obviously has the potential for improving the efficiency of information flow within and between organizations and facilitating coordination between inter-dependent activities [9]. Moreover, where networking has been incorporated in professions where much distributed working¹ takes place, it offers the opportunity of integrating formerly separate stages of the design process and enabling the co-authored objects of production to flow more freely between the different parties involved. For example, networking technology enables the same files to be displayed simultaneously on different terminals, allowing the users connected to the network to be constantly updated of the progress

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1992 ACM 0-89791-543-7/92/0010/0346...\$1.50

of each other's work. Such capabilities also provide the means by which innovative CSCW applications (e.g. shared editing, co-authoring) can be put into practice.

The implementation of networking technology in a work setting is obviously going to have a significant effect on the working practices of an organization. The transition from a set-up of single user systems, to a network of interconnected users and computer systems, will inevitably bring about changes in how tasks are carried out (e.g. synchronously or asynchronously), when activities are performed and who performs them. Moreover, the sharing of files and system components (e.g. plotters) will require additional methods of managing the flow of work through the organization. A key issue, therefore, is how well organizations will adapt to these changes; will it be a smooth transition or will it be fraught with difficulties? If the latter is true, how can the kinds of problems that are likely to arise be identified and characterized such that they can inform the development and implementation of future network systems in organizations?

Chandrasekaran [2] comments on how a contemporary organization is itself a distributed system of subdivisions which, "when successful, mesh together in a miracle of purposefulness, but when the overall structure strays too far from the changing environment, it resembles a maladaptive dinosaur" (p3). Appending another distributed system, in the form of a computer network, to a complex human network already in existence, likewise, could either clog up the arteries or facilitate the flow of blood in that organization. The success of introducing networking for the purpose of integrating inter-independent activities in and across organizations, therefore, will depend largely on supporting a flexible relationship between the two distributed systems. To this end, it is important to examine current patterns of the social organization of inter-dependent activities and how groups of people working together adapt to changing socio-technological set-ups.

The development of distributed computer systems, however, has largely been concerned with specifying the requirements for platforms on which applications for sharing information and communication are to be built [6]. While being concerned with various user issues such as access, privacy and availability of objects there has been very little consideration of the social, cultural and cognitive aspects of changing the way information is transmitted, constructed and manipulated throughout an organization. In particular, there has been little research into the processes of coordination (e.g. control, management, negotiation,

delegation of responsibility and exchange), which are so integral to the successful cohesion of a group.

The aim of this paper is to examine the means by which a group of inter-dependent workers interact with each other and a computer network. In particular, it focuses on the communication process that takes place when breakdowns occur between the interconnected networks of people and technology. An analysis is presented which interprets the events as they unfold in relation to the socio-cognitive issues of shared understanding, transmission of knowledge and distributed problem-solving. From an applied perspective, it considers the need to develop 'coordination software' in combination with 'preventative software' to facilitate the mediation of interactions between users and the computer network.

ANALYSING INTER-INDEPENDENT ACTIVITIES

There is a growing concern in CSCW research that attempts to model the working activities of organizations for the purpose of using them as the basis for developing CSCW systems is misguided. Robinson and Bannon [24], for example, discuss the problems of assuming that a model can somehow capture the objective reality of work organizations. In the introductory chapter of the ECSCW'91 proceedings, Schmidt [26] describes such attempts as "reckless as a daredevil mounting a Bengal tiger" (p1). Simply, the reality of work settings is that they are continuously evolving in response to the ever-changing demands of the situation. Hence, distributed work inevitably involves the vagaries of negotiation and re-negotiation. For example, in an empirical study of how spreadsheets are used in a work setting, Nardi and Miller [22] comment on how their construction and use flow in fluid, informal ways and how cooperation among spreadsheet users has a spontaneous self-directed character.

But if distributed work activity is so variable and complex, how is it possible to carry out any sensible research that can lead to a better understanding of it and which can also inform the design of better or more useful CSCW systems? One approach that is currently receiving considerable interest is to examine the design process itself [20, 24]. Another approach is to observe the actual goings on in a work setting using various ethnographic and microgenetic methods [1, 18]. Whilst not being application-driven, such research can be highly informative insofar as it can provide very detailed accounts of important qualitative structures of activities that are critical to the coordination of distributed work activity, which it is argued could go unnoticed through the use of modelling techniques. This is achieved by examining the intricate details of seemingly, simple phenomena and rendering the invisible and subtle changes in the interactions and actions that take place as explicit.

The recognition of the importance of analysing activities in situ has resulted in a number of studies that have examined the details of the communicative and collaborative processes that take place in work settings, where new technological artefacts have been or are in the process of being introduced [7, 8, 10, 13, 25, 27]. A number of recurring themes have emerged from these studies. In particular, a general observation is the instrumental role played by cultural and technological artefacts in mediating distributed activities. In a study of an airline operations

room, Suchman and Trigg [27] discuss the multiple roles played by both paper documents and on-line computer information displays in coordinating the work of the operators. Rogers [25], also discusses the important social role played by a shared information display, in the form of a whiteboard, to coordinate the activities of a group of engineers working closely together.

Another major finding from the empirically-based studies of work-in-setting is the flexibility of groups to adapt the division of labour to the dynamic exigencies of the environment. Hutchins [14] notes how a ship's navigation team were able to restructure their collective activities in a systematic way in response to an unexpected change in the informational environment, when a piece of the navigational equipment failed. Moreover, the participants solved the problem initially without any conscious reflection, suggesting that the solution was discovered by 'the organization itself' before being learned or discovered by any of the individuals. Similarly, in a study of a London Underground control room, Heath and Luff [10] found that when a crisis arose the controllers were able to manage the difficulties through an emergence of a flexible division of labour. They adapted their working practices to the fluctuating local task needs without any explicit recognition of a new strategy emerging. Specifically, the controllers developed the practice of overhearing each other's conversations in conjunction with overseeing each other's actions, while simultaneously engaged in distinct and apparently unrelated tasks.

A NEW THEORETICAL FRAMEWORK

The studies described above indicate the importance of having technological systems that provide mechanisms which support flexible social interactions and work practices. In particular, they suggest that work settings need to provide a variety of mediating mechanisms that allow the workers to organize and coordinate their working activities in relation to one another. From a theoretical perspective, they have raised a number of questions concerning the adequacy of existing conceptual frameworks to account for the types of organizational change and collaboration which have been observed [14]. In particular, the finding that working groups discover new solutions to unexpected changes in the local environment, largely through implicit and intersubjectively constructed processes, begs for accounts that explain the nature of these processes, how they emerge and continue to adapt. This is especially so, in view of the fact that current organizational [see 12], social and cognitive accounts often have failed to recognize the dynamic and intersubjective processes by which the members of the group organize their work and develop informal codes of practice.

In response to the lack of theoretical treatment of local adaptation to organizational and technological change, several researchers have begun developing new analytic frameworks by utilizing a diversity of theoretical constructs. These include the application of activity theory [7, 16], conversational analysis [10], distributed cognition [14], anthropological theories of culture [19] and organizational/community memory [23].

A main concern of the new framework building is determining the nature of the social interactions and cultural mechanisms by which knowledge about activities

and tasks is communicated in a group setting. Krauss and Fusell [15], for example, point out how the dynamic process of knowledge transfer between different individuals, who share mutual knowledge, beliefs and assumptions, is highly complex, entailing several interwoven levels of intersubjectivity. Drawing from Clark and his colleagues theory of *common ground* [3, 4] they describe how in order for some item of information to be mutually known by A and B, it is necessary but not sufficient that A and B know X. A must also know that B knows X, that B knows A knows X, that B knows that A knows that B knows that A knows X and so forth. Obviously when more than two individuals are involved, the layering of who knows what and how this is transmitted becomes considerably more complex. To ground the process by which knowledge is mutually transmitted, the theory assumes that when communicating, participants try to establish that what has been said has been understood. How this is achieved, however, varies considerably from one situation to the next [5].

The major problem of grounding communication in a shared arena where much knowledge and many beliefs are mutually assumed between the participants, is that numerous misunderstandings can arise. In particular, in collaborative situations where many inter-dependent interactions take place between members of the group, it is possible for the confounding of expectations to occur, having repercussions for subsequent interactions and actions. While one individual might assume that the other/s know what is meant by a particular speech act or action – in terms of the actions which the other person/s should carry out in response – the other/s may interpret the situation quite differently and carry out a different set of actions (or not perform any action) than what the first individual assumed they would do.

Although, in general, it may be the case that the context of the particular social interaction can provide sufficient grounding from which the other individuals of the group can interpret correctly the intentions of each other, when much is left *unsaid* the ongoing situation can often be interpreted quite differently by the various individuals. Hence breakdowns in the coordination of distributed activities can arise, because of a failure of the participants to effectively co-construct a compatible meaning of the particular situation. For example, Rogers [25] found how informally established mechanisms of coordination easily broke down, because the various individuals often had different expectations of which mediating mechanism was currently active.

When a group of people work together, via a computer network, it is highly probable that misunderstandings of the ongoing situation be exacerbated. One reason is because of the changes in the way users interact with the networked technology. In particular, the need to constantly negotiate, monitor and check with each other when sharing files, printers, etc., is highly vulnerable to coordination breakdowns [25]. For example, the occurrence of file corruption, missing data and printer blockages can be attributed to misunderstandings between the users of the network. Moreover, on discovering such problems, attempts by the users to resolve them are likely to be further hampered by misconceptions and faulty transmission of knowledge of the on-going situation.

Besides considering the nature of social interactions that facilitate the coordination of a groups' inter-independent activities, therefore, it is necessary to analyse the extent to which multiple misunderstandings can percolate in collaborative work settings. As the following case-study illustrates, slippages in the shared understanding of which remedial actions are required to solve a breakdown in the computer network, are common and very time-consuming.

THE ETHNOGRAPHIC STUDY: THE CIVIL ENGINEERING PRACTICE

Details concerning the nature of the engineering practice that was observed in this study are described elsewhere [25]. Briefly, the type of work that the civil engineering group carry out, includes producing sets of design plans such as street designs and improvement grading plans. This requires much iterative design, with different engineers working on the same drawing plans at different stages of the design. The introduction of a networking system, comprising of workstations, PC's, file servers and plotters was intended to support this process. Specifically, it has enabled files of the drawing plans to be transferred automatically between the various designers and drafters. The networking software that has been installed also allows the same files to be accessed simultaneously at different locations. This means that the engineers can open up any file at any time, for referencing purposes, but not be able to work on the same file simultaneously, because of the problem of potential data clashes. The introduction of the computer network also has enabled an expensive electrostatic plotter to be shared between the different design groups within the practice.

The engineering company is divided into a number of semi-autonomous design groups; each comprising of a core group, who continuously utilize the networking facilities, and an attached support group consisting of management, administrative staff and non-CAD designers. The core groups usually consist of six engineers (two project managers, two designers and two drafters) who work very closely together. Each core group is located in an open plan working area with a particular network configuration. This is linked up with a central workstation (R7), that is located in another part of the building. It is the most powerful machine and is directly connected to the electrostatic plotter. The local network that was set up for the design group, that was observed in this study, comprised of three powerful workstations, two PC's, a file server and an ink plotter. The various components had been configured so that one of the workstations (R5) acted as the main communication centre, that connected to R7. Thus, to send a file from one of the PC's or workstations to the electrostatic plotter, required transmitting it via R5 and R7.

A problem of printing files through the network, from either the ink plotter or the electrostatic plotter, is that they have to be sent in a suitable format that is acceptable by the different plotters. This requires the creation of plotfiles that specify the correct parameters for the particular kind of plot requested. A blockage in the network can arise if the commands are not correctly put in. Files can also get 'stuck' in the network for no apparent reason.

The detailed analysis presented below, describes how the design group dealt with a network blockage incident that

occurred during a typical morning's working activities. Given that this type of problem is an emergent property of the networking environment, there were no written procedures available which the group could follow to remedy it. Instead they have had to first discover the problem, then identify the possible causes and finally develop ways of dealing with it. Of interest here, is how that process has been established and the extent to which it has been maintained, in the context of the problem of mutual knowledge transmission in an ever-changing environment.

Detailed analysis of distributed troubleshooting

As mentioned above, the plotter problem is caused by various conditions and can be resolved by carrying out certain actions. Through having successively dealt with this type of breakdown over a period of many months, the group have developed a shared understanding of an informal troubleshooting schema (see Figure 1). Primarily, the schema consists of various working hypotheses about the nature of the problem and a set of remedial actions that should be followed on detecting various symptoms. An important difference between the various remedial actions that can be carried out, is the extent to which they disrupt the other processes that are currently active on the network. This can range from having no effect, to disrupting all processes that are connected to the central workstation, R7.

Symptom:	Electrostatic plotter not printing
Default condition:	R7 is blocked
Default cause:	<i>One plotfile is not able to be printed preventing any others being printed</i>
Default action:	Kill the guilty plotfile and resend (localised, no effect on other processes)
	If symptom persists: Redo primary action If still continues to fail then take:
Secondary actions:	i) Publicise the problem and inform group of intended action ii) Reboot R5 (requires anyone in the design group linked to R5 to stop working on process)
Possible cause:	<i>Bad communication between workstations and PC's</i>
Possible cause:	iii) Recreate plotfile <i>Bad element in plotfile</i>
Possible cause:	iv) Ask system manager to reboot R7 (requires every one from all the design groups linked to R7 to stop working on process)
Possible cause:	<i>Overloading of R7</i>

Figure 1. The troubleshooting schema.

What appears on the surface as a fairly straight forward routine, that should be relatively easy to implement and maintain (especially, since the problem has arisen and been dealt with, previously, on numerous occasions), however, has not been established. While the engineers are all familiar with the various aspects of the schema, it appears

that they have immense difficulty in coordinating their activities and expectations of what each other knows about the network problem. Alternatively, 'ontological drift' [24] can arise whereby the various members of the group develop differing interpretations of the state of the problem. Moreover, such discrepancies can often go undetected for long periods of time.

Stage 1. Identification of the plotter problem.

Mark² (a drafter) initially recognizes that there is a problem in the network, having failed to print plotfiles off the electrostatic plotter. He associates it with the default condition of the troubleshooting schema and subsequently adopts the default action of resending the plotfiles (see Figure 1). After resending the plotfile, he brings up the plot monitor queue (qstat³) on his PC monitor. This shows its status, as represented in the plot queue, as having arrived at R7. This should mean that it has been accepted for printing and hence should be coming off the plotter. To determine if the plotfile actually is being printed, Mark goes downstairs to where the electrostatic plotter is located. He discovers that the plotfile has not been printed and so tries again. But, after thirty minutes of repeating several times the cycle of resending plotfiles, checking qstat and going downstairs, it becomes blatantly clear to him that the default action has not worked and that the plotfiles are still getting stuck in the network.

The problem Mark is confronted with is a discrepancy between the physical evidence (the plot is not being printed off) and the information displayed on qstat (stating that the plotfile has arrived and, therefore, should be printing). The mismatch of external representations eventually leads Mark re-consider his model of the likely cause of the problem. This triggers the decision to take a secondary action, which is to reboot R5. However, unlike the default action, which simply removes the guilty plotfile, rebooting will kill everyone else's processes that are connected to R5. Taking such action, therefore, requires Mark to inform the others of his intended remedial action.

Stage 2. Making public the problem.

Mark transmits his discovery of the plotter breakdown by first informing Gary (a designer), who is the only other person sitting in the shared working area at the time (see Figure 2). Without waiting for any acknowledgement from Gary, he continues by stating that he is going to reboot R5. Gary responds in amazement, that Mark is not following the implicitly accepted code of practice which is to *ask* everyone else connected to R5 if it is OK to reboot R5. However, Mark does not concede to Gary's rebuttal. This could be because he considers it unnecessary to inform anyone else as Gary is currently the only other person in the shared working area (and hence connected to R5).

This dialogue indicates a mismatch of expectations between two members of the group as to what constitutes the 'correct' way of transmitting knowledge of how to take control of a shared object. On the one hand, it appears that one interpretation of the situation is that a communicative act is required to inform, while on the other, it is assumed that a negotiative process be set up whereby permission to execute the action is first requested. However, the two participants do not seem to take any notice of each other's interpretation of the situation. Indeed, the ensuing discourse appears more like two parallel streams of talk,

which occasionally interact with each other, and whereby the context of the problem is continuously re-interpreted with respect to each other's concerns (see Figure 2). Mark's stream of talk is largely mediated by the information that is displayed on his monitor, while Gary's contributions are oriented towards prescribing the correct practice.

(10.05)
Mark: We got trouble here. <looks at monitor>
Gary: (laugh) Why's that?
Mark: Sent out. There it's [
Gary: [won't go?
Mark: It won't go. (...) We need to reboot this one.
Gary: Reboot this one?
Mark: Yeh!
Gary: No, no, no, no, no, no, [
Mark: [yes, because
Gary: [yes, because right now all of
us are hooked up. You reboot you kill me.
Mark: (laugh)
Gary: You have to check, make sure no-one else is on
Rick 5.
Mark: No reply. (...) Sent. Trouble at end. It arrives and
then not comes out. <brings up qstat again> (...)
Gary: <moves over to Mark's PC and peers at his monitor>
Check it, make sure what the problem is first.

Figure 2. Making the problem public.

Following the dialogue with Gary, Mark elaborates his working hypothesis to me as to what he perceives the problem to be (see Figure 3). The reason why he has decided that it is necessary to reboot R5 is because he thinks there must be a communication problem between his PC and R5. Even so, he is still not certain what is preventing his plotfiles from reaching R7.

(10.10)
Mark: The plot is stuck and not going. It's somewhere.
Y: So [
Mark: [So we send it over there. It goes over there <points
to R5> and it stay in there and not go down there.
Y: Oh, it's not going downstairs?
Mark: No, it sometime happen like that. It should be
queuing here. <points to plot queue monitor> We cannot
queue it.
Y: So you don't know where the problem is?
Mark: <shakes his head>

Figure 3. Enumeration of working hypothesis.

Stage 3. Wider transmission of the problem.

At this point Phil (the project manager) walks into the shared working area. Mark immediately shouts out to him that the plotter is 'not going'. Phil responds by asking him for how long this has been a problem. Mark answers by saying about half an hour. Phil then asks if Kate (another drafter) is downstairs, to which Mark simply shrugs his shoulders and the conversation ends. There is no explicit communication between the two as to the possible cause. Furthermore, Mark, does not volunteer further information about his working hypothesis of the problem being a communication problem between his PC and R5.

Here, we see the beginning of a mismatch of 'local' knowledge between the two. Instead of trying to co-

construct a mutual understanding of the current state of the problem through learning more about Mark's experience, Phil decides to seek more information from the system manager (Brian), whose job it is to oversee the running of R7. He calls Brian, who is physically located next to R7, and begins to explain to him the problem in the wider context of the design group's activities – the need to get a set of plots off for checking (see Figure 4). Brian replies and Phil's immediate response is confrontational. There is no negotiation between the two about the cause of the problem or the remedial action required at that stage. Instead, Phil ends the conversation abruptly, not having got any closer to resolving the problem.

(10.15)
Phil: Hey Brian, do you know if the plotter's all blocked
up? (...) We're all doing electrostatics for one last check.
We send em [
Brian: <replies over phone>
Phil: Don't you ever get your plot queue thing to roll?
Brian: <replies again over the phone>
Phil: No? Hang on, I'll call back later.

Figure 4. Telephone call between project manager and system manager.

In the meantime Kate has returned from downstairs, after waiting for her plotfiles to come off the plotter. Phil asks her if she has had any success and for an update of her situation. She begins to explain how many plotfiles she has created and what has happened to them but is interrupted in mid-sentence by Phil. Whilst listening to Kate, he has simultaneously brought up qstat on R5. At a glance he has been able to deduce the same information that Kate is verbally communicating to him; that the files have been sent and are showing up as having arrived and therefore should be printing. Here, again is another example of a technological artefact mediating the communication. By 'reading' off the screen, Mark is able to confirm his interpretation of the situation as represented by the network, with Kate's explanation.

At this point all the members of the group collaborate, by sharing their acknowledgement of the problem (see Figure 5). They spend considerable time informing each other of the status of the various plotfiles that have been submitted and compare this with the information displayed on the screen. There is much verbal repetition of the events that have occurred. The outcome allows Phil to become aware of the discrepancy between the network's representation of the plotfile and its actual state, which Mark had experienced earlier but which had not been communicated between the two. The difference between their understanding, however, is that Phil's knowledge of the problem is constituted through a combined interpretation of the system's representation of the plotfiles and his understanding of the information obtained from the other engineers' problems, while Mark and Kate's understanding, also is based on their direct interaction with the network – in terms of sending plotfiles. As a consequence, it appears that Phil's interpretation of the situation is at odds with Mark's and Kate's. Rather than assuming it to be a communication problem between the terminals on the network, Phil initially attributes the disappearance of the plotfiles to the possibility of another person inadvertently walking off with the printed files. Hence, at this stage he is

inclined to believe the network and consider the problem as a human error. Such a possibility is not even considered by Kate or Mark since they 'know' the local situation at the plotter, having made frequent visits downstairs.

(10.25)
Phil: On the computer it just says 'waiting in line', so you got trouble on that computer too?
Kate <nods>
Phil: You got trouble, too?
Gary: About half an hour ago.
Mark: Me too.
(...)

(10.35)
Phil: OK, we're going to re-submit them and see what happens. (...) Someone may have walked off with it or something. (...) Check that one <pointing to filename on screen> (...) it's not that big a file, either. (...) Yes, you see something isn't right. (...) Rick 5. (...) Can't do it. (...) Oh yes, I got a couple stuck here.
Kate: Oh?
Phil: One says routing and the other says queuing. {muttering to himself}. I'm going to reboot this machine. We could try that, try that first and I'll get back to you. {in loud voice} OK everyone off. I'll reboot this one!
Mark: [So it's going?]
Gary: Nothing really good!
Mark: [It's gone away?]
Phil: Everybody off, Gary. Are you on Rick5, there? Anyone?
Gary: Yes! (...) What a mess eh? Can't believe it. <starts to whistle>

Figure 5. Acknowledgement of each other's problems and collaborative troubleshooting.

On finding that the default action initially has not worked, Phil, like Mark, repeats the cycle of killing the plotfiles and resending them. However, each time he repeats the remedial action he requests Kate or Mark to act as relays for him – whereby they go downstairs to see if the plotfiles are being printed. Mark complies, even though he has been through the same process before and has as a result established a different working hypothesis. Neither, does he say anything about his current hypothesis. Consequently, the shifting of the responsibility of the problem from being one individual's to becoming the group's – but with a different person taking up the main role of troubleshooter – causes it to go back a stage in its resolution.

Stage 4. Taking up secondary action.

Having gone through the motions of the first stage of identifying the problem, Phil decides to take up the secondary action of rebooting R5 – which is what Mark had wanted to do in the first place. The manner in which Phil decides to carry out the action is similar to Mark's, but even more abrupt (see Figure 5). Rather than asking if the others linked to R5 mind being booted off the network, he categorically demands that everyone just get off it. Gary, who happens to be connected, responds by shouting out 'yes'. In contrast with the previous conversation with Mark in the same situation, however, he does not try to point out the 'correct' code of practice to Phil. Alternatively, he makes a jocular comment on the whole state of affairs. The

different organizational roles of Mark and Phil result in different discourses. Whereas Mark is a drafter, and hence of lower status than Gary, Phil is of higher status as a manager, and thus cannot be reprimanded in the same manner.

On another level, the discourse between Mark, Phil and Gary is remarkably similar to the previous interaction between Mark and Gary. But this time it appears as if three streams of conversation are taking place, with each participant ignoring the contributions of the others. This is most evident when on hearing Phil mention that one plotfile is routing and another queuing, Mark asks if Phil has succeeded in sending the plotfiles. On not receiving a response from Phil, he repeats the question. But again Phil does not respond and so Mark is not able to determine what Phil meant by his earlier comment. Here again, is another example of a lack of communication preventing the participants from updating each other about the knowledge and understanding of the problem.

Stage 5. The critical shift in working hypothesis.

Whilst Phil is rebooting R5, the others have started working on other tasks. Other engineers from the design group also enter into the working area requesting information and sets of files from Phil. Phil switches between the various demands, while continuing to focus on the network problem. The effect of rebooting R5 should be to unblock the connections between the peripheral terminals and R7. However, the information provided by qstat suggests that this secondary action has also failed (see Figure 6). At this point, Phil is still interpreting the anomaly with respect to his original hypothesis.

(10.40)
Phil: I don't understand. I think if I reboot {mutter}. There should be something down there. (...) S9 was backed up, S8 was just sent down there.
Kate: Yer.
Mark: No Rick5 got some trouble routing. I mean stuck forever. I do it Saturday and it come back Monday, still in there.
Phil: You reboot, yer?
Mark: I reboot it.
Phil: Where?
Mark: Gone {inaudible mumble}.
Phil: You did that on Saturday?
Mark: I did it on Saturday. Yer only one sheet {inaudible as mumbles last part to himself}.
(...)

(10.45)
Y: What do you think the problem is now?
Phil: Oh, a queuing problem, one machine not talking to another. (...) Carry on and then the signal get crushed.
Y: And you loose the file?
Phil: No, just the plotfile. They can recreate them.

Figure 6. Shift of working hypothesis.

There is then a critical turning point when Mark boldly volunteers his secondary working hypothesis that the problem is a communication failure between the peripheral terminals and R5. He backs his case with evidence of having had similar problems when rebooting R5 three days earlier. Such a disclosure enables Phil to consider the

problem in a different context and as a result he shifts his hypothesis to match up with Mark's current working hypothesis. Transmission of knowledge is explicit at this point. There is, however, no explicit confirmation between the two or the others, who are listening in on the conversation, of what the new information that has been disclosed means with respect to what action should be taken next. Instead the dialogue simply tapers off and the engineers continue with their respective activities. After several minutes, Phil resumes the conversation but on a different topic.

Phil reboots R5 three more times. However, each time the same discrepant information appears on qstat. At this point, he changes his current working hypothesis (i.e. a bad connection between the workstations and R5) to the other secondary cause of the troubleshooting schema (i.e. a fault with the elements in the plotfiles). He mutters this to himself twice and the necessary remedial action. Meanwhile, the others are talking with each other about another issue and so do not overhear him. His change in working hypothesis, therefore, is not transmitted.

After rebooting R5, Phil informs the group that they can start working on the network again. Although there has not been any explicit confirmation with the others, there is an implicit expectation by Phil now for the others to adopt the secondary action of recreating plotfiles⁴ instead of resending them. However, the assumption is grounded only through Phil's interactions with the network. Kate 'picks up' the new expectation and, accordingly, recreates her plotfile. On the other hand, Mark continues to follow the remedial action of killing his stuck plotfile and simply resending it. Based on his earlier question to Phil, as to whether the routing problem had 'gone', his not 'picking up' the new course of remedial action, could be because he is still assuming the problem had been due to communication failures at R5 but which he assumed has now been rectified by Phil.

The mismatch in expectations, as to which remedial action should be taken, continues undetected for some time. As a result the routing problem persists and still no plotfiles are printed. After about ten minutes, Mark asks Kate whether she is re-sending to the plotter the plotfile for file S9 again (see Figure 7). Whilst acknowledging Kate's reply that she is now recreating her plotfiles, Mark still does not 'pick up' on the expectation that he, too, is expected to recreate his plotfiles. It could be that Mark interprets Kate's comments to mean that only the plotfile S9 was bad and needed recreating.

(10.55)
 Mark: Kate, are you sending this one downstairs?
 Kate: Not yet. I'm {waiting to see if BOZO⁵ get's stuck}
 Mark: Create BOZO?
 Kate: Yer we have to recreate it.
 Mark: No?
 Kate: Yer!
 Mark: Again?
 Kate: Yer, Phil reckons there was a lot of bad elements in that file.
 Mark: Oh.
 Kate: And wanted us to recreate it.

Figure 7. Explicit communication of new working hypothesis.

Even though Kate has recreated her plotfile for S9, she still finds that it is getting stuck in the network. She points this out to Phil who confirms with her that she has actually sent a recreated plotfile. However, throughout their conversation it does not occur to either of them that Mark might still be re-sending a 'bad' plotfile. Such an omission seems perfectly reasonable, given Kates previous conversation with Mark about Phil's expectation that they adopt the new course of action to recreate plotfiles.

Another ten minutes passes, during which there is no interaction between any of the engineers. Kate returns again, from being downstairs at the plotter. She says aloud that the plotfile for S9 did not go. Phil returns and Mark, Phil and Kate talk in unison. At this point it occurs to Kate that Mark may not actually be following the alternative remedial action. On asking Mark if he has been continuing to re-send his plotfile (see Figure 8), Mark admits to not having done so. Phil, on overhearing this admission, severely reprimands him. Mark, however, is not given an opportunity to justify his actions, even though there may have been a legitimate reason for him not recreating his plotfiles.

(11.05)
 Kate: Is this the same plotfile?
 {Kate, Mark and Phil begin talking at the same time.}
 Mark: We resend the plot?
 Kate: The old one?
 Phil: 9 and 8. (...) You resend the old one?
 Mark: <All heads turned towards Mark> Yeh.
 Phil: No, no, no, no, no, no, no! Shouldn't do that. If it doesn't go, it's bad. Is probably why it was screwing up. You have to recreate it. Did you recreate it?
 Mark: No.
 Phil: Don't do that. <brings up qstat and types in various commands> When you realize something screw up, recreate it.
 (...)
 Kate: Happens every time.
 Gary: Never fail
 Mark: (laugh)
 Gary: Never fail.

Figure 8. Public disclosure and reprimand of Mark's misconduct.

The public reprimanding of Mark, enables the group to explicitly reify the secondary action of needing to recreate a plotfile if a networking problem persists. Hence Mark's admission to not adopting the alternative action of recreating plotfiles, at the undeclared time when Phil had decided that it should be the new procedure, provides grounding for the heuristic to become explicitly established as an appropriate code of conduct.

The four of them continue interacting, with Phil once again taking central control. On his orders, Mark recreates S9 and, likewise, Kate recreates S8. Phil brings up qstat to check whether the plotfiles have been queued for plotting. But still the plotfiles are not being printed off. Phil can not quite believe this and attributes the blockage to the bad elements in S9 again, even though the file has been recreated at this point. Gary also shares this view with him.

There is, though, a crucial difference between the information represented by qstat insofar as the plotfiles are displayed as queuing on the monitor rather than arriving. This should mean that they are actually at the plotter. Phil, however, is still reluctant to change his working hypothesis of S9 being the guilty plotfile – although he does acknowledge the change in the locus of the problem in the network from R5 to R7. Before embarking on another remedial action, however, he requests actual physical evidence of whether the files have been plotted. Once again, Mark is sent downstairs to check the status of the plotter. On returning empty handed, Phil decides that the final measure – rebooting R7 – needs to be taken. He can not do this himself, however, as this is the system manager's responsibility.

Phil calls up the system manager in charge of the running of R7 (who is more senior than the one he talked to earlier) and asks him if he will reboot it. The response is favourable and Phil thanks him profusely. On putting the phone down he announces to everyone that R7 is about to be rebooted. But there is no acknowledgement from any of the others. Instead, they continue working on other tasks not requiring the use of the network. Phil disappears and then returns ten minutes later. He asks Mark to check qstat. Both Mark and Gary bring up qstat and comment that S8 has disappeared from the monitor, S9 is queuing and S10 is still running. The fact that S8 has disappeared from qstat implies that it has finally run off the plotter. However, there are no shared signs of relief. Phil simply asks Mark to go and collect the printed copy of S8 so that he can start the checking process. The production process recommences and the files start to print off in an orderly fashion. No one mentions that the problem is resolved or what is finally accepted as the cause. The incident passes as a new set of demands takes its place.

But the next day the same problem recurs; this time it is Phil who discovers and first makes public that a file he has sent to the plotter has got stuck in the network. And this time, it is Mark who dares to ask him if he has recreated the plotfile. And the cycle of events repeats itself, taking a couple of hours to sort out, but with the engineers playing different roles with a different pattern of misunderstandings, regressions, mutual sharing of knowledge and confirmations.

IMPLICATIONS AND CONCLUSIONS

The analysis has shown how a relatively simple network problem can disrupt distributed work activities quite considerably. Furthermore, even though this type of breakdown has been identified and resolved on numerous other occasions by the group, it has not been possible for a routine solution to be established, akin to following the procedures in the troubleshooting schema outlined in Figure 1. Alternatively, the 'ghosts in the machine' problem emerges in differing contexts, with a different pattern of actions and interactions taking place. Each time, a complex web of communication pathways have to be set up, maintained and coordinated in order to resolve the problem. However, this type of self-organizing appears to be very resistant to being transformed into a crystallized set of working procedures. Instead, the dynamics of the moment-to-moment demands of the multitude of activities taking place in the work setting need to be dealt with in a flexible and reactive manner.

An advantage of this form of local adaptation is that whenever new 'teething' problems arise (which is inevitable given that systems are regularly being reconfigured, with new components being added or replaced and new versions of software being updated), the group are able to modify their ways of working without having to radically restructure their set of working practices. In addition, contrary to the common assumption that increasingly sophisticated technology results in more alienation, the networking technology has brought about more social interaction in the work place⁶. The analysis showed how the engineers need to know what each other is doing on the network. This knowledge is communicated quite freely and often reiterated, enabling all members of the group to be updated of each others progress. However, the transcription also showed, that even though the level of communication may have increased, misunderstandings still frequently occur between members of the group. Critical knowledge which one or more of the group possesses, may fail to be transmitted at the 'window of opportunity' to the other/s engaged in the distributed problem solving. Consequently, a tremendous amount of time is wasted managing the network breakdowns.

Clearly, such a situation is undesirable and counter to the expectation that networking can improve efficiency within and across organizations. The important research issue, therefore, is to determine what types of communicative and informational support are needed to help the group manage networking and communicative problems when they arise. This entails developing resources which can facilitate collaboration and also to seek ways of detecting and recognizing misunderstandings and misinterpretations of situations between members of the group.

As was pointed out in the analysis, a contributing factor to the misunderstandings that arose between the engineers was the discrepant representational states of the files. Specifically, the information provided by the network did not match up with the physical reality of the situation. As a consequence, the engineers had difficulties in diagnosing the cause of the problem, which led to the emergence of differing working hypotheses. An obvious solution to reducing such anomalies would be to develop more informative monitoring systems that reflect more accurately the current status of files in the network. A problem with this solution, though, is the level of 'intelligence' that would need to be programmed into the network, to provide a more transparent interface.

Other kinds of 'coordination' software which could be beneficial include an on-line update of the activities of each of the users and a CSCW-type communication facility, where local negotiations could take place. For example, if a user needed to reboot a communal system it would be useful if they could alert all the other users of their intended action via some form of electronic mail or, perhaps, even set up a negotiating space by which the other users could accept or request that the user wait until they are ready to quit the network.

Another solution could be to provide more 'preventative' software, whereby the system is programmed to have increasing responsibility [see 25]. One potential advantage of delegating the network as manager and organizer, is that

it could prevent some of the undesirable situations arising through misunderstandings or miscommunication. For example, file clashes could be prevented by implementing file-locking mechanisms; printer blockages could be reduced by restricting access to certain file types or users or through automatic rebooting and so on. But in giving the network increasing power and authority, it raises the question of what the resulting relationship will be between the users and the system. In particular, will such a set-up constrain the flexible working practices that exist in an organization to the extent that the communication processes in the group begin to break down? Or will the current trend of facilitating social interaction continue by drawing the users together against the network?

Future research and developments in networking environments will need to consider carefully the nature of the relationships between users, artefacts and mediating mechanisms. The implications of introducing more preventative measures or allowing greater flexibility of user control will need to be analysed with respect to the social functions that come to be served by the network, as more or less responsibility is delegated to it, and how in doing so the network is able (or not) to prescribe certain behaviours in the users [cf. 17]. Conversely, in the social context of their use, set-ups of networks will need to be analysed from the perspective of participants in the pursuits of their users [cf. 23].

ACKNOWLEDGEMENTS

The study was carried out whilst visiting the Cognitive Science Department, UCSD. The research was supported, in part, by grant NCC 2-591 awarded to Donald Norman and Edwin Hutchins from the Ames Research Center of the National Aeronautics and Space Agency in the Aviation Safety/Automation Program.

The paper has developed from the ideas that were being explored in the Distributed Cognition Group led at the time by Ed Hutchins. Thanks to Tim DeWitt of Rick Engineering, San Diego, for allowing me to carry out an ethnographic study of his design group.

REFERENCES

1. Agre, P.E. and Shragar, J. Routine evolution as the microgenetic basis of skill acquisition. In *Proc. of the 12th Annual Conference of the Cognitive Science Society* (1991), LEA: Hillsdale, NJ.
2. Chandrasekaran, B. Natural and social system metaphors for distributed problem solving: Introduction to the issue. *SMC-11*, (1981), 1-5.
3. Clark, H.H. and Marshall, C.E. Definitive reference and mutual knowledge. In *Elements of Discourse Understanding*, A. K. Joshi, I. Sag and B. Webber, Eds., CUP: Cambridge, 1981, pp. 10-63.
4. Clark, H.H. and Shaefer, E.F. Concealing one's meanings from overhearers. *Journal of Memory and Language*, 26, (1987), pp. 209-255.
5. Clark, H.H. and Brennan, S.E. Grounding in communication. In *Perspectives on Socially Shared Cognition*, L.B. Resnick, J.M. Levine and S.D. Teasley, Eds, American Psychological Association: Washington DC, 1991, pp. 127-149.
6. Dollimore, J., Miranda, E. and Xu, W. The design of a system for distributing shared objects. *The Computer Journal*, 34, (1991), 514-527.
7. Engeström, Y. *Learning, Working and Imaging: Twelve Studies in Activity Theory*. Orienta-Konsultit Oy: Helsinki, 1990.
8. Goodwin, C. and Goodwin, M.J. Formulating planes: seeing as a situated activity. To appear in *Cognition and Communication at Work*, D. Middleton, and Y. Engeström, Eds., Sage: London (in press).
9. Hart, P. and Estrin, D. Computer integration: a co-requirement for effective inter-organization computer network implementation. In *CSCW'90 Proc.* (7th-10th October, 1990, Los Angeles), ACM: New York, 131-142.
10. Heath, C. and Luff, P. Collaborative activity and technological design: task coordination in London underground control rooms. In *Proc. of the Second European Conference on Computer-Supported Cooperative Work: ECSCW'91*, L. Bannon, M. Robinson, and K. Schmidt, Eds, Kluwer: Dordrecht (September 25th-27th, 1991, Amsterdam), pp. 65-80.
11. Herbert, A. Editorial: Distributed Systems. *The Computer Journal*, 1991, 34, 481.
12. Hewitt, B., Gilbert, N., Jirkota, M. and Wilbur, S. *Theories of Multi Party Interaction*. TMPI End of Year Report. British Telecom, Queen Mary and Westfield Colleges and the University of Surrey, 1990.
13. Hutchins, E. The technology of team navigation. In *Intellectual Teamwork: Social and Technical Bases of Collaborative Work*, J. Galegher, R.E., Kraut and C. Egido, Eds, Hillsdale, NJ: LEA., 1990, pp. 191-220.
14. Hutchins, E. Organizing work by adaptation. *Organizational Science*, 2, (1991), 14-39.
15. Krauss, R.M. and Fussell, S.R. Constructing shared communicative environments. In *Perspectives on Socially Shared Cognition*, L.B. Resnick, J.M. Levine and S.D. Teasley, Eds, American Psychological Association: Washington DC, 1991, pp. 172-200.
16. Kuutii, K. The concept of activity as a basic unit of analysis for CSCW research. In *Proc. of the Second European Conference on Computer-Supported Cooperative Work: ECSCW'91*, L. Bannon, M. Robinson, and K. Schmidt, Eds, Kluwer: Dordrecht, 1991, pp. 249-264.
17. Latour, B. Where are the missing masses? Sociology of a few mundane artifacts. To appear in *Shaping Technology-Building Society*, W. Bijker, and J. Law, Eds, MIT Press, (in press).

18. Lave, J. *Cognition in Practice*, CUP: Cambridge, 1988.
19. Levine, J.M. and Moreland, R.L. Culture and socialization in work groups. In *Perspectives on Socially Shared Cognition*, L.B. Resnick, J.M. Levine and S.D. Teasley, Eds, American Psychological Association: Washington DC, 1991, pp. 257-279.
20. MacLean, A. Bellotti, V., Young, R. and Moran, T. Reaching through analogy: design rationale perspective on roles of analogy. In *CHI'91 Proc.*, (April 27th-May 2nd, 1991), S.P. Robertson, G.M. Olson and J.S. Olson, Eds, Addison Wesley: New York, pp. 167-172.
21. Middleton, D. and Edwards D. (Eds), *Collective Remembering*. Sage: London, 1990.
22. Nardi, B.A. and Miller, J.R. An ethnographic study of distributed problem solving in spreadsheet development. In *CSCW'90 Proc.* (7th-10th October, 1990, Los Angeles), ACM: New York, pp.197-208.
23. Orr, J.E. (1990). Sharing knowledge, celebrating identity: community memory in a service culture. In *Collective Remembering*, D. Middleton and D. Edwards, Eds, Sage: London, 1990, pp. 169-189.
24. Robinson, M. and Bannon, L. Questioning representations. In *Proc. of the Second European Conference on Computer-Supported Cooperative Work: ECSCW'91*, L. Bannon, M. Robinson, and K. Schmidt, Eds, Kluwer: Dordrecht, (September 25th-27th, 1991, Amsterdam), pp. 219-234.
25. Rogers, Y. Mediating mechanisms in distributed work activities. *CSCW* (accepted for publication).
26. Schmidt, K. Riding a tiger, or computer supported cooperative work. In *Proc. of the Second European Conference on Computer-Supported Cooperative Work: ECSCW'91*, L. Bannon, M. Robinson, and K. Schmidt, Eds, Kluwer: Dordrecht, (September 25th-27th, 1991, Amsterdam) pp.1-16.
27. Suchman, L.A. and Trigg, R.H., Understanding practice: video as a medium for reflection and design. In *Design at Work*, J. Greenbaum and M. Kyng, Eds, LEA: Hillsdale, NJ, 1991, pp.65-90.

iterative process. The sense of distributed refers to the notion that the activities are inter-dependent; an individual or joint activity has consequences for, or is contingent upon another activity, which is currently or will be carried out by another individual or sub-set of individuals in the group. For example, a design may be initially created by one person which is then passed on to other people to review and make their comments. The document is then returned to the first person or is given to someone else to make the proposed changes and so on.

²All names have been changed.

³ The qstat provides information on the status of files sent over the network; they can be either on route, queuing, waiting, on hold or arriving.

⁴Recreating a plotfile requires setting a whole set of formatting parameters which takes much more time and effort than resending a plotfile.

⁵BOZO is the name given when first creating a plotfile.

⁶This was pointed out to me by Geoff Cooper.

APPENDIX: NOTATION USED IN TRANSCRIPTIONS

- [current speaker's talk overlapped by the talk of another
- (...) untimed pause
- <> activity
- { } comment on activity
- () transcriber's uncertainty over verbatim

NOTES

¹Distributed activities are defined as activities which are carried out individually or jointly in a group, as part of an