

Analysing Asynchronous Collaboration.

Patrick G. T. Healey and Nick Bryan-Kinns*

Department of Computer Science, Queen Mary and Westfield College,
University of London, Mile End, London. E1 4NS, UK.

EEmail: ph@dcs.qmw.ac.uk, N.Bryan-Kinns@mdx.ac.uk

* Contact address: School of Computing Science, Middlesex University,
Bounds Green Road, London N11 2NQ, U.K.

In this paper we propose a framework for modelling asynchronous collaboration. Drawing on arguments that a task-based approach is inappropriate for the analysis collaborative work we model collaboration in terms of the artefacts, users, and information flow involved. We illustrate the framework's potential through studies of shared care of diabetic patients, and discuss the implications of this approach for the analysis of collaborative work.

Keywords: Evaluation, CSCW, Communication, Collaboration.

1 Background

Collaboration is an essential part of human activity and the predominant mode of work activity. Computer support for such activity is a relatively recent development in human terms, but a relatively old development in computer science terms. However, the development of effective collaborative systems is frustrated by a paucity of analysis and evaluation techniques that deal effectively with collaborative activity. Although there is now a large body of empirical work on the complex, and often subtle, factors that influence the organisation of collaborative work little progress has been made with techniques or methods that provide for tractable, systematic analyses of a kind that could inform design.

Throughout the literature there are many examples of problems with the use of collaboration support systems. Categories of problems include lack of acceptance (*e.g.* Markus' studies (1990) of why group work products fail to be accepted by users), poor understanding of the purpose of the such systems by users (*e.g.* Orlikowski's studies (1992) of people's perceptions of what constitutes a groupware system), and usability problems with particular aspects of systems.

Such problems have led to numerous studies of collaborative systems. For example, Rouncefield *et al.*'s study (1994) of a small office considering the greater introduction of I.T., and Grønbaek *et al.*'s study (1992) of a Danish engineering company, produce similar forms of descriptions by studying work practices. However, as their understanding and descriptions of both the work and the environment are anecdotal and not based on theories or explanations, their findings are difficult to use generally.

Work such as Bower's study (1994) of the introduction of CSCW tools into an organisation of the U.K.'s central government provide analyses which are inferred from observations of the introduction and use of systems. This analysis highlights the amount of effort that users of groupware systems need to expend to make the systems and their collaborations effective. Bower's study also uses these explanations to propose design solutions to the problems such as: more sophisticated awareness mechanisms, more sophisticated access mechanisms, and more support for management of the CSCW network.

Conversely, van der Veer *et al.*'s Groupware Task Analysis framework (GTA) (van der Veer *et al.*, 1996) supports the modelling of people's work in terms of group tasks, the organisational context, and the artefacts involved. Their framework can be used to develop models of group activities, but there is no explicit relationship to models of systems being used. This means that although group work can be described there is no systematic means of identifying problems, nor of explaining why these problems arise, nor designing to support groupwork.

Bannon *et al.*'s discussion (1997) of *common information spaces* provides a framework for understanding the properties of information objects that cross between different communities of practice. Such understanding can be used to highlight the amount of effort that is needed to bring information from one community into a shared arena. This effort typically involves packaging and filtering the information to make it comprehensible to the intended audience. Their work is similar to that of Bower's study discussed previously which considers the amount of effort required to make collaborations effective. However, Bannon *et al.*'s approach is more powerful as they provide a framework which can be applied to different situations as opposed to Bower's anecdotal analyses.

An example of work in which theoretically based explanations are produced is Roger's work (1992) in which she examines the problems of distributed problem solving. In particular she focused on how problems within the troubleshooting arose, and explained these using an *ad-hoc* framework for representing shared understanding between group members. The framework was also used to explain how shared understandings were developed, and how the problems above were addressed. In this way it did not only describe the situations and problems, but also to explain how the problems arose, and even to provide some insight into how the problems could be resolved.

Inevitably, the investigation of situations in which people collaborate raises the problem of multiple perspectives on the structure and organisation of the joint activity. Participants in a given setting will frequently differ in their view of what actions are involved, how they are integrated and who is responsible for doing what. The management of these differences in perspective are a basic aspect of collaborative activity. For example, Strauss (1988) highlights how co-operative work depends on 'articulation work', a process of negotiating task alignment within a 'working division of labour'. Similarly, Hughes *et. al.* (1992) observe that it is a feature of collaborative work that: "the separation, individuation and combination of activities is accomplished in an accountable way through a collectively developed, negotiated and evolving knowledge and practice" (p.117). On this view, the processes of interaction and negotiation which sustain the working division of labour are constitutive of co-operative work.

These phenomena present obvious difficulties for techniques which use task or goal decompositions as the basis of their analyses of collaborative work (*cf.* Watts

and Monk, 1998). One of the criticisms most consistently directed at 'conventional' HCI has been that it analyses work in individualistic terms; concentrating on an isolated user engaged in some autonomous task (*e.g.*, Schmidt & Bannon, 1992; Heath & Luff, 1992; Jirotko, Gilbert and Luff, 1992). However, the problem raised above is not with individualistic approaches per se, it creates difficulties for any level of analysis (individual, group or higher) at which a privileged decomposition into independent processes, goals or tasks is used. The problem is twofold; not only do different participants in the work take different views on its decomposition, task organisation is continually changing in response to local contingencies and negotiation (*e.g.*, Hughes, Randall and Shapiro, 1992; Bowers, Button and Sharrock 1995; Randall, Rouncefield and Hughes 1995).

In addition to problems with task decomposition, a second, related, consequence of multiple perspectives is that they also engender different views on the significance and meaning of the various artefacts associated with a task. A typical example is provided by Symon, Long and Ellis (1996) who carried out a detailed analysis of the use of a radiography request form (RF) in a hospital. This work demonstrates, amongst other things, how the different individuals who are involved in processing the RF, including nurses, consultants, house officers, radiographers, clerical staff and secretarial staff all interpret the RF and its function in different, sometimes conflicting, ways. Indeed, Star has suggested that a key factor in the success of a collaborative artefact may be the extent to which its design actually sustains multiple interpretations (*e.g.*, Star and Griesemer, 1989).

These difficulties suggest that collaborative activity may require analyses that do not employ notions of task or goal as basic units of analysis and, partly as a result, do not privilege a particular understanding of the objects and artefacts involved. The empirical studies suggest that a promising alternative would be to make communicative processes the central focus of analytic interest. One approach which has developed along these lines is the application of the grounding model for the analysis of interaction with or through computational media (Brennan, 1998; Clark and Brennan 1991). To date, this model has only been applied to the analysis of local, synchronous interaction. We develop this approach and propose a framework for analysing and comparing the organisation of a broad range of collaborative activities which, we believe, captures those aspects which are most significant for the design of collaborative systems. The rest of this paper is organised as follows. We first outline our approach and follow it with an illustrative example. We then give two further contrastive examples which provide further indications of the framework's use. Implications are then discussed along with the use of the framework to inform design. Finally the paper is summarised.

2 Approach

Our approach takes the artefacts and individuals in a setting as the basic units of analysis. Any object that can be used as a document of collaborative activities is understood as an artefact. This includes familiar examples such as a letter, memo, image, database entry as well as allowing for unconventional cases in which objects such as chairs may be used to assess the current state of the collaborative activity (see Anderson, 1996). No attempt is made to determine what a memo or letter specifically means to the various participants in a particular setting, *i.e.*, what it is 'really' a document of. Particular artefacts may be used as documents for a

variety of activities, we discriminate artefacts only on formal/physical grounds and they are considered to change only where, for example, they have been annotated, amended or destroyed.

Direct, synchronous interaction is treated in two ways. Firstly, it may itself be documented. For example, a consultation between a GP and a patient will be documented by a number of artefacts. Typically, there will be an appointment in one or more diaries, a list of actions taken on the patient record, a prescription or referral letter and so on. The second feature of synchronous interaction for our purposes is that it provides an opportunity for individuals to align their interpretations of an artefact, whatever those interpretations may be, through grounding. Grounding is the basic process used to update *common ground* in communication and was originally designed to account for how individuals develop a common understanding of utterances or turns during a conversational exchange (Clark and Schaefer, 1989). We extend this approach to cover asynchronous interaction by treating artefacts as, in effect, superordinate 'turns' that are also grounded between individuals.

The principal advantage of thinking in terms of grounding is that it allows us to exploit distinctions in the *state*, or *level of grounding* in which one of the partners may be in with respect to the shared artefacts involved. Clark and Brennan (1991) define four states of grounding with respect to conversation, and Brennan (1998) defines eight states with respect to grounding with interactive help systems. We draw on these to define five states of grounding with respect to artefacts:

- State 0:** Participant (P) is unaware that artefact (A) exists.
- State 1:** P is aware that A exists.
- State 2:** P recognises A as being of a particular type.
- State 3:** P understands the content of A.
- State 4:** P understands what actions are associated with A.

It should be noted that in the original conversational model, an utterance is understood as successfully grounded at some level only if *positive* evidence has been obtained, at that level, for the mutual belief that the utterance has been grounded. In the current context this criterion cannot, in general, be satisfied. For collaborations that are distributed widely in space and time there is no reliable means, for analysts or participants, of determining if the right mutual belief has been established. We therefore adopt a weaker criterion and assume that individual beliefs about the level of common ground with respect to an artefact are accurate unless there is evidence to the contrary *i.e.* unless we have specific evidence of two or more individuals maintaining discrepant beliefs about each others common ground with respect to some artefact. We initially propose individual grounding levels from observations and records of use *e.g.* whether an individual has seen an artefact (level 1), or whether, based on their roles and experience, we might assume that they understand the content of the artefact (level 3) and associated actions (level 4).

We use this analysis to generate two diagrammatic representations of the informational state of a particular collaborative activity (illustrated in later examples). The first is a Grounding Matrix which summarises the level of grounding for each individual with respect to each relevant artefact. With one exception, this matrix also represents, by default, the mutually assumed level of

common ground for a given artefact for each pair of individuals. If the levels of grounding are different for two participants they are assumed to hold the mutual-belief that their common ground is at the lower level. The exception is where, for some individual and some artefact, the level of grounding is zero. In this case, although their collaborators are assumed to know that they are unaware of the artefact, the individual in question has, of course, no beliefs about it or their collaborators' knowledge of it. These defaults are chosen on the assumption that people are normally aware of which artefacts are associated with a particular collaborative activity and are aware of what level of understanding or involvement each of their colleagues have with each artefact. Where there is evidence that the beliefs are in fact discrepant, we take this to be significant and represent it directly within the Matrix. If, for example, participant A believes B's level of grounding with respect to an artefact is higher or lower than it actually is this is subscripted on A's level of grounding.

To address the specific patterns of information flow a graph representation is used (for example, see figure 1) which addresses the interactions used to support collaboration, and the different levels of grounding that result from those interactions. Individuals, represented by circles, and artefacts, represented by rectangles, form the nodes of this graph. The arcs between them represent the interactions. Asynchronous interactions are represented as thin lines with filled arrow heads and synchronous interactions as thicker lines with open arrow heads. The artefacts associated with each interaction are linked to that arc by a single line which is labelled to indicate the level of grounding achieved for that artefact through the interaction.

3 Illustrative Example of the Approach

A typical example of problems with support for group work is presented here to illustrate the use of the evaluation framework. This, and later examples, are taken from studies of shared diabetic patient care (Kindberg *et al.*, 1999). Health care for diabetic patients is typically organised in terms of three sectors: primary (GPs, opticians, chiropodists, district nurses), secondary (specialists), and self-help. These sectors relate to different degrees of expertise and specialisation in the treatment they provide. Diabetes is a chronic disease; the aim of the care is to prevent the disease and related problems from worsening. As such the self-help level of care is vital. Patients must ensure that they take, and possibly amend, their treatment every day, that they monitor the levels of sugar in their blood, and that they eat appropriate foods.

The primary sector is the patient's first line of contact. They may visit their GP at regular intervals, and typically visit the optician regularly as diabetes can lead to degeneration of eyesight. If the patient's condition is beyond the treatment of the GP then they are *referred* to a specialist in the secondary sector. The referral acts as a co-ordinating device; clinicians instigate future activities through referrals and repeat appointments.

One result of the disparate nature of the shared care is that patients are often treated by several clinicians concurrently (in separate sessions). This concurrent working is *loosely-coupled*; the members of the group work semi-autonomously, and over a long period of time towards the group goal of reducing the degeneration of the patient's health. In the example extracts that follow, a nurse arranges an

appointment for a patient at an eye clinic just prior to the patient seeing the GP. However, when the patient sees the GP the GP, despite looking through the shared notes and database, is apparently unaware that the appointment has been made until the patient produces the appointment letter. The following extract from the transcript relates to the nurse deciding to refer the patient to the eye clinic.

Nurse: Have you got another appointment at the L. [eye clinic]?
Patient: No, I'm under Dr. X. [the Dr. the patient sees next] and nurses here.
N: Right... in that case, we need to make an appointment for the next time the camera is available... 6 weeks.
P: I didn't like the L. - they had wrong address for me and blamed me for not turning up
P: I was under Dr. whatever-her-name-was in G [a clinic].
Nurse looks at paper notes.
N: Yes
The patient then proceeds to air more grievances about the L. whilst the nurse types at the computer.
N: You need to make an appointment.

At this point the nurse has written a letter to arrange an appointment for the patient at the eye clinic and noted the arrangement in the shared database. This letter is printed out at the end of the session and handed to the patient for them to arrange the appointment.

The incident of interest occurs when the patient sees the GP as illustrated in the following transcript extract.

Patient: I hung on [for an eye examination] because I had it done at the L.... I explained to Y. [the nurse] that I can't remember when.
The GP (Dr. X.) then looks through both the paper and computer based patient notes
Dr. X.: Has Y. given you opticians things?
P: Gave me this letter to tell me to go 'round there.
The patient then shows the letter prepared by the nurse to the GP.
X: That's fine.
P: Mr. M [a consultant]. Y said go in a month's time.
X: He'll send us the result.

Although the nurse has noted that they wrote an appointment letter for the patient and informed them when to visit the eye clinic the GP appears to have missed the information. In fact, the GP spends some time looking through the computer and paper based notes (presumably in an attempt to find out what the nurse has done with respect to arranging for the patient to visit the eye clinic) before asking the patient what the nurse has done (Symon, et. al. 1996 make similar observations in the context of radiography). It is through this interaction that common ground about the appointment with the eye examination is indirectly established for the nurse and the GP.

4 Use of the Framework in Analysing Shared Care

In this example there are two main artefacts of interest: the appointment letter, and the shared database entry. There are three participants: the nurse, the GP, and the patient. From the transcripts, video recordings, and observations carried out at the time, the Grounding Matrix shown in table 1 is derived (the numbers refer to levels of grounding as discussed previously).

	Nurse (N)	Patient (P)	GP (GP)
Appoint. letter	4 _{GP:1}	3	4
Database entry	4 _{GP: 4}	0	1

Table 1: Grounding Matrix regarding patient's eye appointment

This illustrates that two artefacts could be used as documents of the patient's eye appointment; the letter and the database entry. Levels of grounding for these artefacts are uneven. The nurse, by virtue of authoring the letter and database entry is assumed to have the highest level of grounding for both. The GP by contrast, has a high level of grounding for the letter but not for the database entry. He reads and confirms the appropriateness of the letter even though it is actually intended for the eye clinic and has only contingently become available to the GP. The organisational expectation is that it is the procedure of making a database entry which should provide the GP with a document of the appointment. We note this by annotating the nurses expectation that the GP will use the database but will only be aware of the letter. However, for whatever reason, the GP has not grounded this information by reading the database and notes. The patient clearly understands the type and meaning of the letter but we assume is not fully aware of the activities it initiates for the various clinicians involved.

The Grounding Matrix highlights one reason why the letter is a more critical artefact in this case than the database entry; it is, on average, more highly grounded for the individuals involved. However, this is also a deviation from the expected pattern of co-ordination. To account for this we need to consider the interactions involved, illustrated in figure 1. In this case the nurse and the GP interact asynchronously through the database whereas the patient interacts synchronously with both the nurse and the GP during face to face interactions. Grounding of participants with respect to the artefacts is indicated by the numbers either side of lines connecting interactions to artefacts. For example, the nurse is grounded to level 4 with the database entry whereas the GP has only reached level 1.

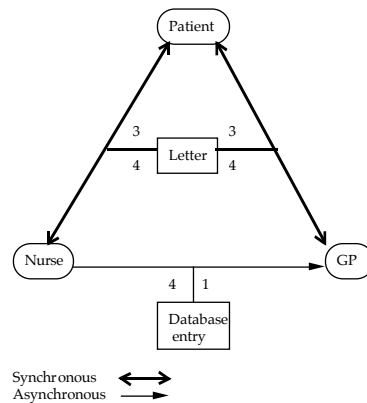


Figure 1: Information flows for patient's eye appointment

This captures the fact that, although the database entry provides a direct informational link between the Nurse and GP, it is the indirect channel provided by the letter that underwrites successful co-ordination in this case. It is effective, we suggest, because there are opportunities for the participants to ground the letter through direct interaction that are not available for the database entry. It may even be that the GP intentionally exploits this fact by preferring to rely on the letter in this case.

5 Further Examples from Studies of Shared Care

Two further illustrative examples are presented here. The first concerns problems with results from visits to other sites, the second illustrates a situation in which support for collaboration is well exploited.

A typical example of problems concerning co-ordination between sites occurs when patients attend eye clinics. The purpose of a visit to the eye clinic is for the eye specialist to examine the patient's eyes and determine what treatment is necessary *e.g.* cataract removal. The tests performed by the eye specialists inform their diagnoses. However, the results of such tests are also of interest to the referring clinician as indicators of the patient's diabetic condition. Typically such results and information about what the eye specialist intends to do for the patient are not returned to the referring clinician. The following extract of transcript highlights one such situation. First the nurse does not know whether the patient has attended an eye clinic and so has to expend effort extracting this information from the patient (after failing to find any relevant information). Second, at the end of the extract the nurse states that they will have to 'follow that up' which clearly will involve the nurse in extra work. Note that although they are interested in the same artefact (test results), they use them as documents of different things -surgical need or diabetic state.

Nurse: How are your eyes? Have you had them checked?

Patient: I didn't 'cos I had them done, but I didn't know when I had them done.

At this point the nurse picks up the paper notes – any information from eye clinics will be in the form of a letter.

P: Do they photograph them?

N: Yes
P: I had them done at the L.
N: They never told us.
The nurse looks through the notes whilst talking to the patient.
N: No correspondence.
N: At the L.?
P: Yes.
N: Need to follow that up.
At this point the nurse stops looking through the notes.
N: Have you got another appointment?

In this example the Grounding Matrix (Table 2) consists of: the patient, the nurse, the eye specialist; the photograph taken, and the results letter. As before, the eye specialist, as author of the relevant artefacts is assumed to have the highest level of grounding. Note that although we have no evidence to demonstrate that the letter exists, we equally have no reason to believe it doesn't and analyse the example accordingly.

	Nurse	Eye specialist	Patient
Photograph	1	4	2
Results Letter	0	4	0

Table 2: Grounding Matrix regarding the patient's eye test

Figure 2 depicts our analysis of the flow of information in this case. The visit to the eye specialist generated two artefacts, a photo of the patient's retina and test results. The expected asynchronous interaction between the nurse and the specialist failed, there is no letter containing test results in the patient record. In response, the nurse uses the synchronous interaction with the patient to establish that a visit to the eye specialist took place. The patient knows that a photograph has been taken (level 2), but does not necessarily understand the meaning of the photo. The nurse, on the other hand, has only reached the first state of grounding with respect to the photograph; the only evidence they have of its existence is from the patient's report. As the patient does not know about the results letter and the nurse/GP have not received one the nurse has no direct evidence that it exists (level 0).

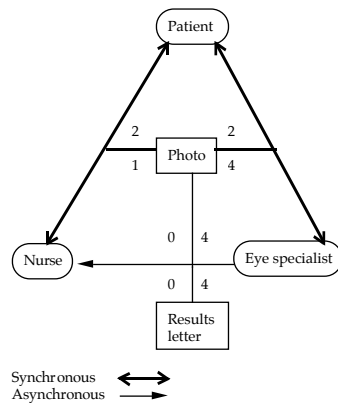


Figure 2: Information flow regarding the patient's eye test

The interest of this example lies in the way in which the failure of the asynchronous communication was both detected and dealt with. The nurse uses the interaction with the patient to ground the existence of one artefact, the photo, as a means of determining whether an examination took place and, in turn, to assess whether another artefact, the test results, could be obtained. A variety of factors may be involved in this situation. One possible factor in the failure to feed back the results is that the specialists, who use them for their own purposes, may not appreciate their potential value to the clinicians. Problems of this kind are pervasive in collaboration between individuals with different fields of expertise and are very difficult to address at the level of system design. Also, conflicts of interest and concerns about confidentiality dictate that the simplest solution -making every artefact available to everyone involved in patient care cannot succeed (*cf.* Symon et. al. 1996). What this and the preceding example highlight is how people can cope with these situations through a strategic and indirect use of artefacts and interactions. An effective system or procedure for the support of this co-ordination would not only attend to the results letter but to understanding which other artefacts could usefully be made available. The analysis framework developed here assists in identifying which artefacts are of greatest relevance to this by attending directly to the levels of grounding and patterns of information flow in each case.

Of course, communication does not always fail in shared care (if it did it could hardly be referred to as *shared* care). This final example concerns a situation in which communication is supported by the systems in place. In this example the nurse attempts to find out about the patient's diet; whether it is appropriate for the dietary requirements of attempting to keep diabetes under control. The nurse has a difficult time finding out the relevant information from the patient and so raises concerns with the GP regarding the patient's diet *via* the shared database using the following comment. From the studies it is clear that such comments serve to both indicate the nurse's opinion about the situation, and as a form of informal request for the Dr. to pursue the theme with the patient.

I found it difficult to ascertain what he was eating etc.

When the Dr. sees the patient (shortly after they have seen the nurse) they first read information contained in the shared database. From this they pick up on the

nurse's comments regarding diet and start the session with the patient by discussing their diet. The gravity of the Dr.'s concern about the patient's dietary situation is reflected by the amount of time they spend questioning the patient in this initial part of the session, and the fact that they return to the dietary theme two further times in the session.

Essentially the comment helps to establish the Dr.'s understanding of the current state of the shared care; that the patient needs to be further questioned about their diet, and that the patient's diet needs to be carefully monitored in the future (not just by the Dr.). In terms of the evaluation framework the communication between the Nurse and the GP about the patient's diet is successful; both have grounded to level 4 as illustrated in table 3 and figure 3.

	Nurse	Patient	GP
Database entry	4	0	4

Table 3: Grounding Matrix for successful information flow




Figure 3: Successful information flow

6 Implications

Having illustrated the application of the framework in specific cases we discuss the broader implications of this approach. As the examples show, collaboration of the kind involved in shared diabetic patient care is predominantly asynchronous and typically takes place *via* artefacts such as letters. Such loosely-coupled work contrasts with the work situations that have most frequently been studied. Schmidt and Bannon note the “... the underlying assumption in most of the CSCW oriented research thus far that the co-operative work arrangement to be supported by a computer artefact is a small, stable, egalitarian, homogenous and harmonious ensemble of people.” (1992, p.15, see also Plowman, Rogers, and Ramage, 1995).

In effect, the present framework abstracts over particular episodes of interaction in order to characterise the broader patterns of communication which underwrite co-ordination. By doing so it has the potential to address a wider variety of different forms of collaboration, not just tightly-coupled communication. We do not assume that more communication, or higher levels of grounding, are necessarily better or even desirable (*cf.* Harper and Carter, 1994). Rather, this framework offers a means

of systematically characterising the pattern and depth of communication in a particular setting and for comparing these patterns between settings. Its potential value consists in the means it offers for assessing the relative importance of various artefacts and individuals to the current work organisation. The key individual(s) in a particular situation can be identified by inspecting the grounding levels across all artefacts for the highest overall score. Similarly, the centrality of particular artefacts can be assessed by looking at the distribution of grounding levels across a group. We believe this framework also clarifies the situations in which workarounds occur and the factors influencing the structure of those workarounds. For example, the preceding analyses highlighted the strategic value to clinicians of exploiting the patient's knowledge of appointment letters or photographs as documents of procedure rather than the artefacts, such as shared databases or results letters that are ostensibly designed for that purpose. This kind of understanding may be of particular value for requirements analysis by facilitating the identification of critical artefacts and interactions. The framework also facilitates higher-level generalisations about patterns of co-ordination. The Grounding Matrices provide a simple visual representation that indicates the extent to which communication is, for example, fragmented, centralised or diffuse in a given setting. Further, it provides a clear indication of which combinations of participants constitute important coalitions or groups in carrying out the work. We believe it is a strength of this approach that we do not assume, say, an organisational or role-based model of group structure, rather it emerges directly from the pattern of grounding and interaction.

The effectiveness of this framework depends to a significant extent on the ease and accuracy with which the Grounding Matrices and information flow diagrams can be constructed. The kinds of highly detailed analysis which have proved successful for understanding closely-coupled co-ordination are not practicable for more distributed, largely asynchronous, interaction. An advantage of the current framework is that it is designed to apply to situations in which only partial data is available and does not demand the use of full transcripts or other types of more sophisticated microanalysis. A variety of less intensive techniques, *e.g.*, interview, tracer studies, document audit are available which can establish the critical parameters of the framework; typical patterns of interaction, the set of relevant individuals and artefacts, and the pattern of grounding. A deeper question is whether Grounding Matrices and flow diagrams capture the factors that are of greatest relevance to the conduct of collaborative work. We have argued that task-based analyses are inappropriate in these cases because they impose a particular understanding of the work organisation on a setting. The validity of our alternative approach is currently being assessed in the design and evaluation of system for shared patient care.

7 Design of a Shared Artefact

Using examples like those discussed above we have developed an understanding of the nature and requirements of collaborative work in shared patient care. On the basis of our analysis we have designed a system which aims to engender more effective collaboration through the promotion of higher levels of grounding. This system, referred to as mPathy, is illustrated in a screenshot in figure 4 (see Kindberg *et al.* (1999) for more detail). For each patient, it provides a shared

overview of the history of interactions between the relevant clinicians and other health workers associated with that patient. Each clinician can use the workspace to view who has been active with the patient, and assess the state of the communication between them. The clinicians involved in the patient's care are represented by icons on the left with documents they have added to the workspace arranged from left to right on a time line associated with each clinician. In this example four clinicians are visible: an optician, diabetes sister, GP (C. Day), and specialist (S. Good). The workspace promotes grounding to level 1 by making all documents in the workspace visible to all the individuals in that workspace. To promote grounding to level 2 (recognising type), several types of icon are used to distinguish between e.g., letters between clinicians, reviews, or consultation. In the simple case, all documents in the timeline can be read by anyone in the workspace, modified distribution is indicated in two ways. Firstly, a specific intended addressee is indicated by an arrow between participants, for example, on 14 Feb. 1997 C. Day sent a letter to the specialist to request a review. Secondly, where, for reasons such as confidentiality, distribution is restricted a padlock appears on the timeline of each individual who is not currently able to read the document. For example, C. Day's consultation and review of the patient's case are not currently available to the optician. Grounding to level 3 is promoted by providing tick marks to indicate when a document has been read by an individual.

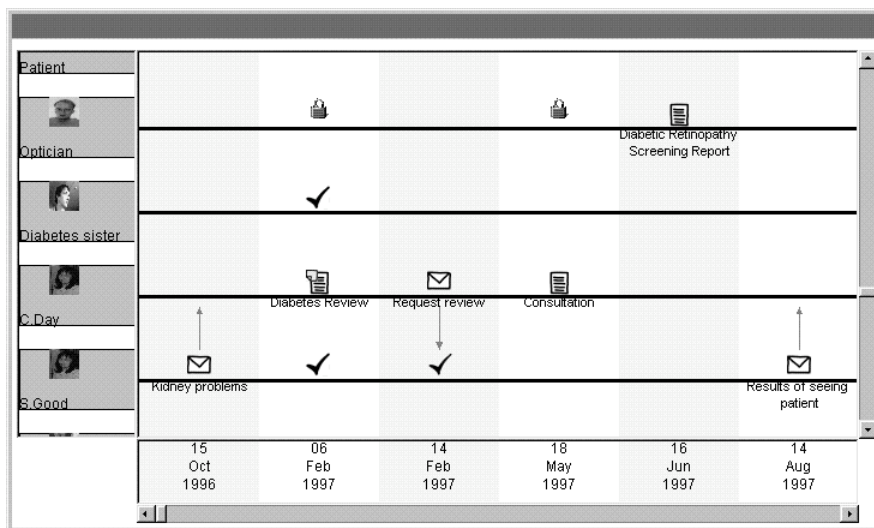


Figure 4: Screenshot of mPathy

The design of the system thus aims to promote both an awareness amongst the clinicians of what patient-related artefacts there are and to underpin their common ground about the current state of patient care.

8 Summary

This paper presented a framework for modelling collaboration in terms of artefacts, information flow, and participants' levels of grounding. This approach contrasts with other work which attempts to model the task and goal structures involved in

group work. By concentrating on levels of grounding, our framework abstracts at a level that avoids the problems associated with the fluid nature of tasks and task allocation and which undermine task based approaches. Furthermore, our Grounding Matrices and diagrams of information flow provide simple visual representations which can be employed to highlight and interpret the nature of the collaboration.

Examples of loosely coupled collaboration from studies of clinicians involved in shared care for diabetic patients were used to illustrate the analytic utility of the approach. Similarly, we proposed how such analyses of collaboration could be used to inform design of systems. Our future work involves the further development of the framework for the analysis of other forms of asynchronous collaboration and its use as a basis for making comparative judgements about the adequacy and effectiveness of collaborative systems.

Acknowledgements

The work presented in this paper is part of the Mushroom project funded by UK EPSRC grants GR/L14602 (1996-97) and GR/L64300 (1998-2000). We would like to thank all those involved in the project, the clinicians who have given us their valuable time, and the patients who agreed for us to observe them.

References

- Anderson, R. (1996). "A Security Policy Model for Clinical Information Systems." *Proc. IEEE Symposium on Security and Privacy*.
- Bannon, L., & Bødker, S. (1997). "Constructing Common Information Spaces." *Proceedings of ECSCW '97*, pp. 81-96.
- Bowers, J. (1994). "The Work to Make a Network Work: Studying CSCW in Action." *Proceedings of CSCW '94*, pp. 287-298.
- Bowers, J., Button, G. and Sharrock, W. (1995) "Workflow from Within and Without: Tehcnology and Cooperative Work on the Print Industry Shopfloor". In Marmolin, H., Sunblad, Y. and Schmidt, K. (eds.) *Proceedings of the Fourth European Conference on Computer-Suported Cooperative Work*, September 10-14th, Stockholm, Sweden. pp.51-66.
- Brennan, S. 1998. The Grounding Problem in Conversations with and Through Computers. in Susan R. Fussell, and Roger J. Kreuz eds. *Social and cognitive approaches to interpersonal communication*. 201-225. Mahwah : Lawrence Erlbaum Associates.
- Button, G. (1992) "The Curious Case of the Vanishing Technology" In G. Button (ed.) *Technology in Working Order: Studies of Work Interaction and Technology*. pp. 10-30. London: Routledge.
- Clark, H.H. and Brennan, S.E. 1991. Grounding in Communication. 127-149 In Resnick, L.B., Levine, J and Behrend, S.D. (Eds.) *Perspectives on Socially Shared Cognition*. Washington DC.: American Psychological Association.

- Clark, H.H. and Schaefer, E.F. 1989. Contributing to discourse. *Cognitive Science*, 13: 259--294.
- Grønbaek, K., Kyng, M., & Mogensen, P. (1992). "CSCW Challenges in Large-Scale Technical Projects - A Case Study." *Proceedings of CSCW '92*, pp. 338-345.
- Harper, R. and Carter, K. 1994. Keeping people apart: A research note. *Computer Supported Cooperative Work* , 2, pp.199-207.
- Heath, C. and Luff, P.(1992) "Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Control Rooms." *Computer Supported Cooperative Work* , 1. pp.69-94.
- Hughes, J. A., Randall, D. and Shapiro, D. (1993) "From Ethnographic Record to System Design: Some experiences from the field." *Computer Supported Cooperative Work* ,1, pp.123-141.
- Jirotko, M., Gilbert, N. and Luff, P. (1992) "On the Social Organisation of Organisations" *Computer Supported Cooperative Work* ,1, pp.95-118.
- Kindberg, T., Bryan-Kinns, N., and Makwana, R. (1999) "Supporting the Shared Care of Diabetic Patients." *Proceedings of ACM GROUP '99*, Phoenix Arizona
- Markus, M. L. (1990). "Why CSCW Applications Fail: Problems in the Adoption of Interdependent Work Tools." *Proceedings of CSCW '90*, pp. 371-380.
- Orlikowski, W. J. (1992). "Learning from Notes: Organisational Issues in Groupware Implementation." *Proceedings of CSCW '92*, pp. 362-369.
- Plowman, L., Rogers, Y. and Ramage, M. (1995) "What are Workplace Studies for?" in Marmolin, H., Sunblad, Y. and Schmidt, K. (eds.) *Proceedings of the Fourth European Conference on Computer-Supported Cooperative Work*, September 10-14th, Stockholm, Sweden. pp.309-324.
- Randall, D., Rouncefield, M. and Hughes, J. A. (1995) "Chalk and Cheese: BPR and ethnomethodologically informed ethnography in CSCW" in Marmolin, H., Sunblad, Y. and Schmidt, K. (eds.) *Proceedings of the Fourth European Conference on Computer-Supported Cooperative Work*, September 10-14th, Stockholm, Sweden. pp.325-340.
- Rogers, Y. (1992). "Ghosts in the Network: Distributed Troubleshooting in a Shared Working Environment." *Proceedings of CSCW '92*, pp. 346-355.
- Rouncefield, M., Hughes, J. A., Rodden, T., & Viller, S. (1994). "Working with "Constant Interruption": CSCW and the Small Office." *Proceedings of CSCW '94*, pp. 275-286.

- Schmidt, K. and Bannon, L. (1992) "Taking CSCW Seriously: Supporting Articulation Work". *Computer Supported Cooperative Work (CSCW) 1*: pp. 7-40.
- Star, S.L. and Greisemer, J. R. (1989) Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology. *Studies of Social Science*, 19, pp.387-420.
- Strauss, A. (1985) "Work and the Division of Labour". *The Sociological Quarterly*, 26 (1): pp.1-19.
- Symon, G., Long, K. and Ellis, J. (1996) "The Coordination of Work Activities: Cooperation and Conflict in a Hospital Context." *Computer Supported Cooperative Work* v.5 n.1 p.1-31.
- van der Veer, G. C., Lenting, B. F., & Bergevoet, B. A. J. (1996). "GTA: Groupware Task Analysis - Modelling Complexity." *Acta Psychologica* 91, pp. 297-322.
- Watts, L. and Monk, A. (1998) "Reasoning about tasks, activities and technology to support collaboration." *Ergonomics*, 41(11), pp. 1583-1606.