Cues to Mutual Knowledge

N. Bryan-Kinns, P. G. T. Healey, D. Papworth, and A. Vaduuva
IMC Group, Department of Computer Science, Queen Mary, University of London
Mile End, London. E1 4NS
nickbk@dcs.qmul.ac.uk

Abstract. We show that asynchronous collaboration can be made more effective by providing cues to mutual knowledge. We demonstrate this by empirically comparing two user interfaces used to support collaborative work. Our position is that effective collaboration is characterized by more co-ordinated and speculative interaction, and that cues to mutual knowledge help participants develop common ground for interaction. We also suggest that more effective collaboration is indicated by increased reliance on expectations of others' knowledge which is characterized by implicit references to shared documents and ideas.

Introduction

Collaboration; the stuff that happens between people when they work together. But how do we understand collaborative activity and design for it? There is a plethora of research on understanding collaboration from approaches which decompose the cognitive structures of collaboration (Johnson and Hyde, 2003) to task-agnostic work which focuses on the nature of the communicative media utilized in collaboration (Watts and Monk, 1998). We can analyze collaborative activities as distributed cognitive systems (Hutchins and Klausen, 1996) and use this to understand how information is shared and transformed in the system. We could think of collaborative activities as activity systems transforming objects in a work context (Issroff and Scanlon, 2002) in order to understand the conflicts inherent within a system. The content of communication could be analyzed (Olsen et al., 1993) to tell us whether we focus more on the technological issues than the actual work we are attempting to undertake. All these approaches, and more, shed light on the nature of our interaction with other, and with the systems we use to support us. They can be used to direct our design decisions, and to allow us to evaluate the systems we build. For instance, a range of key attributes of systems that support us in collaborating with others who are not in the same space have been developed over the last twenty years or so. These emphasize the importance of features such as shared and consistent representations (Robertson, 1997), and awareness mechanisms (Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002). By designing representations which are shared and consistent between remote spaces we reflect the nature of co-located collaboration where we share the same aural, and to some extent visual space. Similarly, awareness mechanisms attempt to support ongoing awareness of others' activities even when we are not co-located. Such awareness helps to co-ordinate the collaborative activity and, in remote collaboration, is typically supported through representations of the current activity of others, and indicators of past contributions. Such representations are central not only to work oriented collaboration, but also to support creative collaborations (Bryan-Kinns and Healey, 2007). Gutwin and Greenberg's approach focuses on real-time aspects of workspace awareness, in particular, the who, what, and where questions. For instance, who is present in the shared workspace at the moment, what are they doing, and where are they looking at the moment. Designing user interfaces that allow participants to answer these questions gives collaborators an awareness of what is going on in the group on a moment by moment basis. However, our understanding of group work is not just informed by what is going on at one moment, but also by what has happened in the past. In particular, we rely on presumptions about who knows what about what has gone on, and beliefs about what we think other people know about what we know (cf. Clark, 1996). The question then becomes one of how to support the development and sustenance of *mutual knowledge* in collaborations – the set of beliefs individuals have about others and their beliefs – that enables communication and collaboration to progress without continuous affirmation and reaffirmation of understanding.

Common Information or Mutual-Knowledge?

The distinction between information that is shared, in the sense of 'held in common' and information that is mutually-known is illustrated by the Conway paradox (see Barwise, 1989). Consider two people, Ann and Bob, playing cards. Each has an ace. They each know, amongst other things, that 'at least one of us has an ace'. This is shared information in the sense that they both know the same thing. Now, if another person, Claire say, asks them "Do you know anything about other's cards?" they will answer "no". Moreover, they will still answer "no" if Claire asks the same question a second or third time.

Consider what happens if Claire now tells Ann and Bob that "at least one of you has an ace". What was common -but independently known- information is now mutually-known i.e., Ann and Bob both know that the other knows that "at least one of you has an ace”. From each individual's point of view very little has changed. Claire has only told them something they each already knew. If Claire now asks, as before, “Do you know anything about other's cards?” they will again answer "no". However, this first response now gives them the additional information that each of them has an ace. They each now know that the other's "no" entails that the other has an ace (since, if they didn't they would be able to answer "yes" to the first question).

One of the difficulties in modeling mutual-knowledge is that it involves this problematic form of self-reference - my knowledge of your knowledge involves your knowledge of my knowledge and so on. Clark Marshall (1981) adopted what is known as the ‘shared environment’ response to this problem. Instead of securing mutual-knowledge we use the cues available to us as a basis on which mutual-knowledge could reasonably be assumed. The simplest ground for such mutual-belief is physical co-presence. If I can see a cup between us, and I can see that you can see it too, then we can (defeasibly) assume that we both know there is a cup between us. Likewise if someone says "at least one of you has an ace" to us we can, all things being equal, assume that we mutually-believe that that at least one of us has an ace.

The grounding model, developed by Clark and co-workers (e.g. Clark, 1996) explores the processes through which people provide one-another with evidence for establishing the layers of mutual-belief necessary for effective communication. Various levels can be distinguished. For example, we might both know that something was said but not what was said. Or we might both know what was said but not what it meant. For example, Clark and Brennan (1991) defined four distinct states of grounding with respect to an utterance:

State 0: B didn't notice that A uttered any u.

State 1: B noticed that A uttered some u (but wasn't in state 2).

State 2: B correctly heard u (but wasn't in state 3).

State 3: B understood what A meant by u.

The central focus of the grounding model is understanding how people manipulate the 'shared environment' to achieve different levels of mutual belief (Clark, 1996). In most of these analyses the focus is on synchronous conversational interactions where people can provide each other with particularly direct forms of linguistic and paralinguistic evidence that they understand each other. There are difficulties in directly applying notions of common ground to the design and evaluation of synchronous collaboration (Koschmann and LeBaron, 2003), but we believe that it can nonetheless be used in a productive way in design.

Brennan (1998) exploited the grounding model to design system feedback that provides cues to the current level of grounding that has been reached with respect to the user's goals. Healey and Bryan-Kinns (2000) extended this approach to modeling the role of artifacts in supporting mutual-knowledge in asynchronous collaboration.

In this paper we report on an experimental exploration of the impact, on collaboration, of cues that are designed to help people maintain mutual-beliefs about the current state of that collaboration. We start our journey with a description of the experiment itself. We then move on to hear the results of our experiment and draw these into discussion. Our journey ends with the conclusion in which we set out the plans for further explorations in the domain of support for remote collaboration.

The Experiment

In order to investigate the effects of providing cues to grounding state on the effectiveness of collaboration in a shared workspace we distinguish three classes of cues that could be simply graphically represented in a computer based interface:

- *First Order*: Cues to the activities of an individual in an environment. For example; icons indicating whether an email has been read, forwarded or replied to.
- *Second Order*: Cues to the activities of others in an environment. For example, the 'read receipt' indication that someone has received and opened an email.
- *Third Order*: Cues that support mutual-beliefs about people's activities in an environment. For example, a conversation about an email that everyone received.

Hypothesis

If the maintenance of mutual, rather that individual, beliefs about the current state of a joint activity plays an important role in collaboration then, we predict, third order cues should have a positive impact on it. In particular we would predict that increased support for mutual-belief should lead to:

- *Less conservative contributions* – more activity related communication, and more discussion than in ineffective collaboration. For us, communication in effective collaboration focuses on the activity at hand, rather than the technological or co-ordination problems that need to be resolved in order to collaborate.
- *Co-ordinated use of artifacts* – participants share ownership of artifacts and manipulate each others' artifacts. This moves beyond reading others' contributions as it entails explicitly adding to, or referring to, each others' contributions. Such activity relies on a shared understanding of the public events that have occurred so far, and an understanding of what is important to the current state of the joint activity, and an understanding of what is meant by the content of the artifacts. All of these rely on the existence of common ground between participants.
- *Less reliance on explicit references to artifacts* – rather than referring to artifacts explicitly, there is an increased reliance on assumptions about mutual

knowledge i.e. assumptions about others' knowledge of the existence of artifacts in the workspace, their content, and their meaning.

Materials

For this experiment two versions of a shared workspace application were developed: Npathy and Mpathy. They were designed to be functionally equivalent, and to differ in the cues they provide about the pattern and state of communication, or grounding, amongst the users of the shared workspace as follows:

- Npathy: 1st and 2nd order cues
- Mpathy: 1st, 2nd, and 3rd order cues

The shared workspaces Npathy and Mpathy were developed in Mushroom (Kindberg et al., 1996) – a CSCW architecture that supports the development of applications for collaborative work based around a notion of shared workspaces. In Mushroom shared objects embody both client and server functionality and are replicated in a 'persistence' domain. In the experimental setting, the Npathy and Mpathy workspaces were individual clients per subject, which had a view onto the relevant part of the persistence domain for their subject group's data. Two user manuals were produced, one for each version of the system.

Npathy

The Npathy workspace provides a title bar at the top, a shared workspace area for documents, a command bar at the bottom and a list of users (referred to as the user menu) on the left hand side as illustrated in figures 1a and 1b. The user menu consists of a strip of icons of other workspace users with their name and, where available, a thumbnail picture of them. These icons change colour to indicate who is currently active in the workspace during a given session. Documents are immutable[1] and represented by different icons according to their type (memo, test, image, document). Each document also has three text fields displayed below the icon showing creation date, author and subject. Small icons were also added to documents as iconic 'superscripts' to indicate if a document has an attachment or annotation. Because annotations could be added at any point, colour was used (red) to signal if they were new.

Although all workspace users share the same set of documents and people, their view of the workspace varies. Firstly, Npathy allows individual users to arrange documents in the workspace according to their preferences by dragging the

1 The Npathy and Mpathy applications were actually versions of an application developed for the support of Diabetic patient care (Kindberg et al., 1999). A key design constraint inherited from the medical application domain was that the body of a medical record or other document should not be altered once introduced into the workspace. Nonetheless, effective support for collaboration requires that comments and notes can be made on each document. In Npathy and Mpathy these functions were supported by functions for adding marginal notes or annotations to the existing documents or creating a new document that cross referenced another through an attachment mechanism.

document icons across the workspace (all documents are always visible in the workspace). They can also choose to rearrange the order of the icons on the user menu. In addition, each user receives information about their pattern of activity in the workspace. A tick icon is added to each document in the workspace that the user has read during a current or previous session. Furthermore, where the author of a document has elected to deny a user access to that document this is indicated to them through a padlock icon attached to the document.

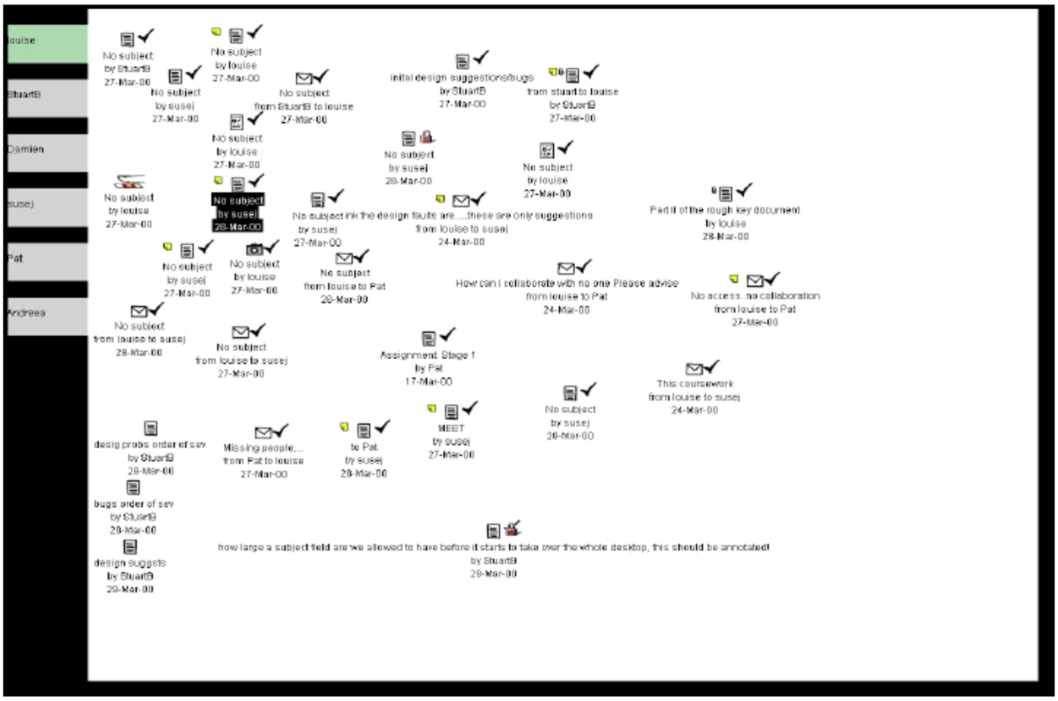

Figure 1a: Screenshot of Npathy

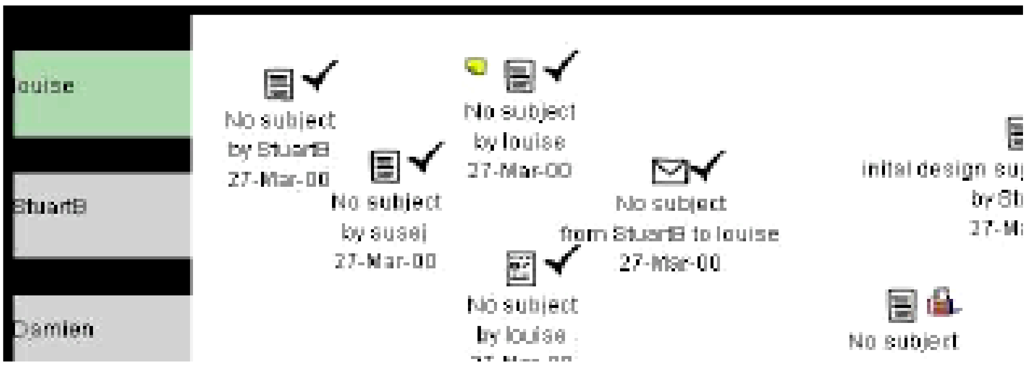

Figure 1b: Detail of Npathy

Mpathy

The Mpathy application reproduces all the document information available in Npathy and additionally provides extra cues about the pattern of collaborative activity with respect to the workspace. As illustrated in figure 2, Mpathy adopted a more structured approach to the representation of the document workspace. Instead of allowing users to individually determine the position of documents we adopted a time based matrix representation which reflects the model developed in Healey and Bryan-Kinns (2000) - in Mpathy, a timeline is associated with each user. As each user contributes new documents they are ordered along the timeline according to the dates on which they were introduced into the workspace.

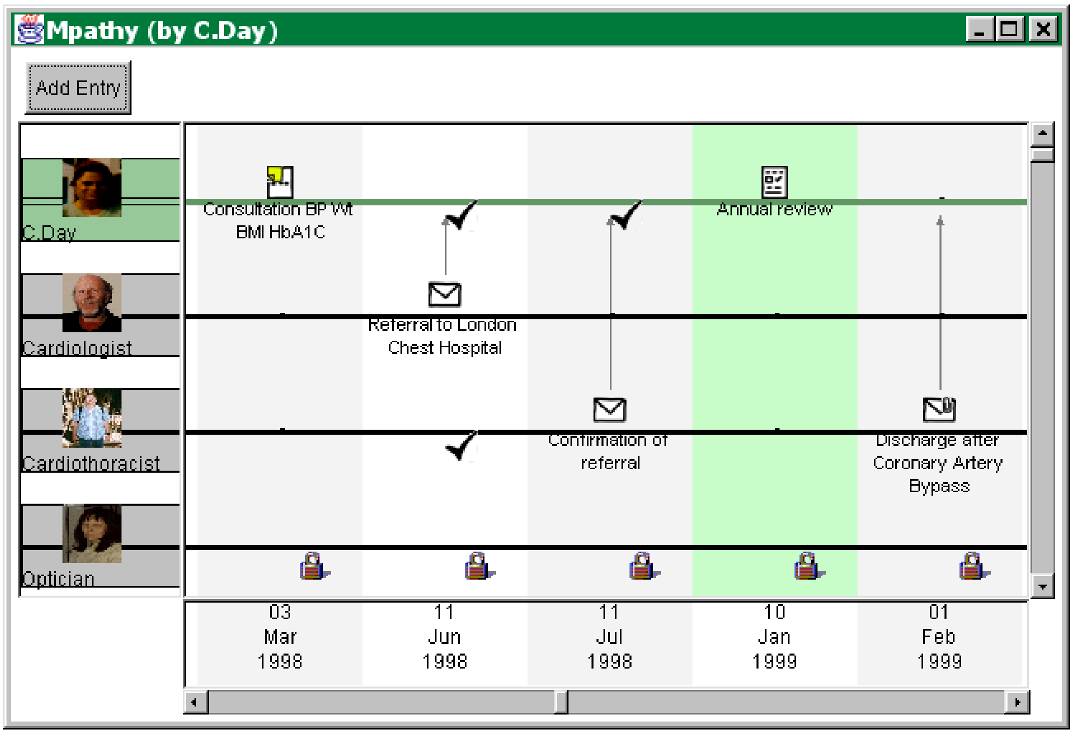

Figure 2: Mpathy user interface

The document workspace in Mpathy provides all users with information about the level of grounding within the group for each document or artifact. Firstly, while Npathy only shows if a user themselves has read a document, Mpathy also shows whether each other member of the workspace has read the document by displaying ticks on the corresponding part of their timeline. For instance, in figure 2, both *C Day* and the *Cardiothoracist* have read the *Referral letter* sent by the *Cardiologist* on *11 Jun 1998*. Secondly, while Npathy only shows a user if they are denied access to a document, Mpathy also shows all other users' level of access to a document. For example, in figure 2 we see that the *Optician* does not have access to any of the documents displayed. Thirdly, in Npathy, the intended recipient of a memo can only be determined by reading the memo whereas in Mpathy the intended recipient of a memo is directly indicated to all users. For example, the recipient of the *Referral letter* sent by the *Cardiologist* on *11 Jun 1998* is illustrated by the grey arrow pointing to *C Day* in figure 2.

Aside from the addition of cues to grounding the other major difference between Mpathy and Npathy is that in Npathy all document icons are visible to the user, whereas in Mpathy users have to scroll to find document icons. This difference, and its possible impact on user behaviour, is returned to in the discussion section.

Subjects

Thirty subjects were recruited from an MSc. class in Computer Supported Co-operative Work. They carried out the evaluation as part of a coursework. They were randomly assigned into 10 groups (5 for Mpathy, and 5 for Npathy) of 3 with a single workspace per group. In order to introduce some role asymmetries and promote subjects' use of access control on documents all groups also included three of the authors. One author was present to provide user help, another provided technical support in the event of crashes or bugs and the third set and marked the assignment.

Procedure

User manuals for Mpathy and Npathy were distributed to all the subjects. Subjects were instructed that they should collaborate, using the workspace, to produce 3 documents:

- A list of design problems ranked according to their severity.
- A list of design suggestions ranked according to their potential to improve the effectiveness of the workspace.
- A list of 'bugs' ranked according to their severity.

It was emphasized that as far as possible all assignment related collaboration should take place in the workspace. Subjects were informed that the assessment of the coursework would be based only on the documents in the workspace, and the amount of activity in the workspace. They were given 14 days to complete the assignment.

Experimental Design

The experiment was carried out with a between subjects factors of user interface (Mpathy/ Npathy). We assumed that each individual in the experiment provides independent data which constrains the dependent measures we can use. The dependent measures derived from our characterization of effective collaboration outlined at the start of this section were:

- *Number of contributions* (documents and annotations). Increased contribution of documents indicates to us an increased amount of co-ordination between participants and an increase in the willingness to communicate.
- *Average size of documents*. In terms of efficiency of collaboration, the size of the documents is related to the number of contributions. A small number of long contributions indicates a conservative attitude to collaboration which is less co-ordinated than situations in which there are a large number of short contributions. With large numbers of short contributions participants indicate a willingness to interact with others and to share the work space. Moreover, they rely on the user interface and shared knowledge to help them navigate the shared documents. Anecdotally we suggest that large numbers of small documents makes the collaboration more akin to conversation than email or letter writing.
- *Number of times documents were read*. Unlike the number of contributions, we argue that increases in the number of times documents are re-read indicates less efficient collaboration. For us, increased document reading would indicate reduced knowledge about the content of each document is i.e. people repeatedly read documents to remind themselves of the content. When people increasingly read others' documents it indicates a willingness to collaborate, but also a lack of shared knowledge about the content of the shared documents.
- *Number of cross-references between documents*. Creating cross-references between documents implies an understanding of the content of both documents that are linked. Where the documents are created by different people this indicates effective collaboration which relies on mutual knowledge about the content of both documents.

In addition, dependent measures of the topics of textual content of the documents were developed drawing on previous analyses of shared document creation and editing (Olson et al., 1993). We developed three categories of document content topic:

(1) *References to other documents* – where subjects refer to the content of other documents or previous discussion. This is divided into whether subjects referred to documents explicitly e.g. "*part of my 'discussion' suggestion in my 27/03/2000 document*", or implicitly e.g. "*I agree with most of what J. said*". From our position, indirect references indicate more effective communication as there is more reliance on assumptions about others' knowledge of the content and meaning of documents i.e. there is more shared knowledge about the public events so far.

(2) *Requests for action* – where subjects request action from others e.g. "*I suggest we all put onto the workspace our ideas and then take it from there*". High numbers of requests for action indicate to us uncoordinated

activity where subjects have to explicitly co-ordinate their action rather than relying on assumptions about shared goals and plans.

(3) *System related issues* – where subjects discuss technical difficulties with the system or test out its features e.g. "*just testing out the memo feature*". A high proportion of system related discussion would indicate that the design of the interface is interfering with the interaction. This is essentially a group measure, and not related to the provision of cues to mutual knowledge per se, but provides us with an indication of whether there are system related issues confounding our results.

Results

One group from the Mpathy condition failed to carry out the assignment and they were dropped from the analysis. A criterion level of 0.05 was adopted for all statistical tests. To preserve statistical power we analyse throughout by individual rather than group. Although in each case our measures are logically independent it is possible that, for example, the fact that one group contains a particularly active participant might increase the activity of each other participant in that group. Fully addressing this problem would require a much larger sample size. Here we note that this increases the risk of a Type I error or 'false positive' in the results reported below.

Logs of activity on the system were collected for a 45 hour period prior to the deadline. This period was chosen both because the subjects would have become more experienced with using the system by this point and the approaching deadline meant that higher levels of activity would occur. The global statistics for this period show that there were a total of 2000 object accesses of which 57% were Read accesses, 34% Modifications and 8% Creations.

The number of documents or annotations created by each individual in the two conditions was calculated. This was entered into a one-way analysis of variance with application type (MPathy vs. Npathy) as a single between subjects factor. This showed a reliable main effect of interface ($F(1,25)=6.01$, $p=0.02$) with subjects in using the Npathy workspace making an average of 5.1 contributions each and subjects using the Mpathy workspace contributing an average of 11.1 documents each.

The number of times each individual made a read access to any document in the workspace was calculated. This was analyzed in an analysis of variance with workspace type as a between subjects factor. This showed no reliable difference in average number of read accesses for users of Mpathy (59%) or Npathy (57%) ($F(1,25)=2.47$, $p=0.12$)

The average size of documents (number of lines of text in the document) created by participants in the two conditions was calculated. The showed no reliable difference in average size of documents for participants of Mpathy (6.88) or Npathy (8.61) ($F(1,25)=0.89$, $p=0.37$).

The average number of cross-references made between documents was calculated which showed no reliable difference between Mpathy (4) and Npathy (1.2) ($F(1,25)=4.26$, $p=0.07$).

Table 1 and figure 3 illustrate the results of analyses of document content topics in terms of the previously detailed categories: (1) Explicit or implicit references to other documents, (2) Requests for action, and (3) System related topics. Table 1 shows the average number of occurrences of each topic followed by the percentage of the overall identified topics. Figure 3 shows the average occurrences of topics as a percentage of overall topics identified, and the variance of occurrences between groups. None of these showed a significant difference in document content between Npathy and Mpathy.

Condition	**Direct**	**Indirect**	**Request**	**System**
Npathy	64 (41%)	15 (11%)	34 (21%)	43 (27%)
Mpathy	52 (34%)	20 (13%)	37 (22%)	33 (31%)

Table 1: Average numbers of topics of document content per condition

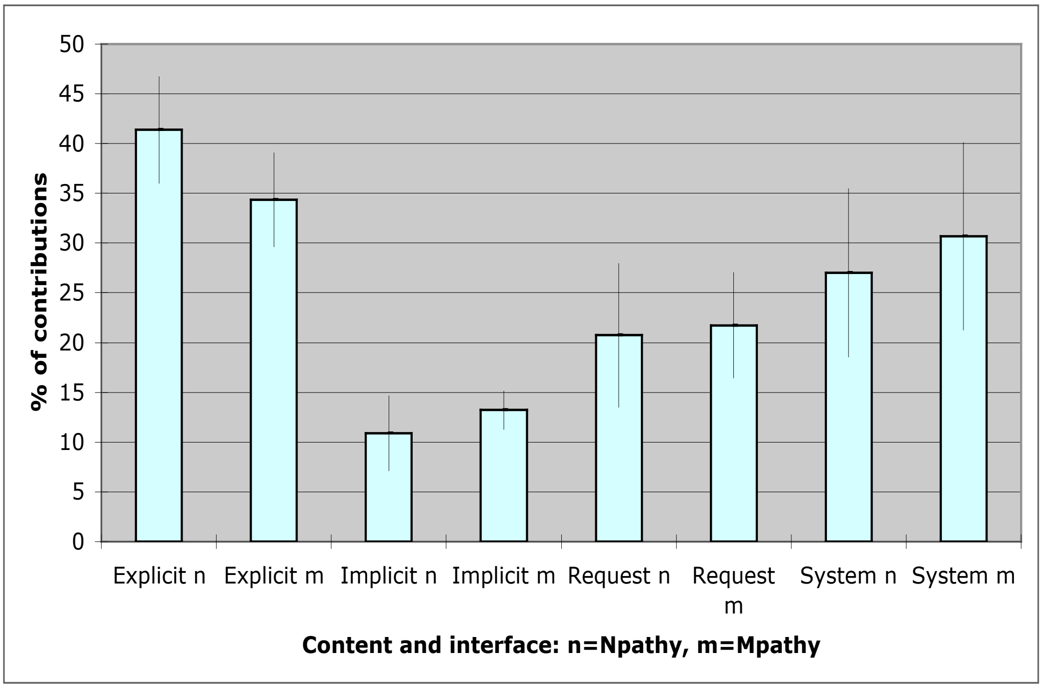

Figure 3: Average percentage and variance of topics of contributions in different conditions

Discussion

This paper examines the effect of cues to mutual knowledge on the effectiveness of collaboration. We compared one user interface which provided 3 levels of cues to mutual knowledge to one which only provided 2 levels. Consistent with our expectations, we found that the groups who were provided with cues to mutual knowledge did indeed collaborate more effectively as characterised by increased contributions, increased co-ordinated use of documents, and relatively more indirect references to mutual knowledge. These indicators are discussed in the following paragraphs, but, given the nature of the two user interfaces examined, caution must be exercised before making assumptions about the critical differences between the two user interfaces. We designed the interfaces to provide different cues to possible grounding with reference to artifacts, in doing so, we necessarily created interfaces which were different in several respects, not just which cues to grounding were visible. For instance, participants could see all document icons in Npathy, but not Mpathy, participants could exercise control over where icons were placed in Npathy compared to Mpathy, and Mpathy made time a prominent aspect of the interface. We discuss these differences later in this section and argue that whilst any of these could have caused the differences we found, the overall results are consistent with our expectations that providing cues to mutual knowledge increases the effectiveness of collaboration.

Before launching into a discussion of the results, we would like to highlight that we found that both groups of participants performed the same amount of reading of documents, as indicated by there being no significant difference between the average number of read accesses by participants ($F(1,25) = 2.47$, $p=0.12$). This indicates that although Npathy users contributed fewer documents, they had to read each document more often which may be because it either contained several points of information (which may have been kept in individual documents in Mpathy), or it was harder to identify which documents were pertinent. Our analysis of the size of documents indicates that there were probably not more points per document in Npathy as the documents were, on average, similar in size to those in Mpathy. Either way, it points to more inefficient collaboration in Npathy users as documents have to be read each document more often than in Mpathy.

More contributions

In the experiment Mpathy users contributed almost twice as many documents as Npathy users (Mpathy average: 11.1 contributions, Npathy average: 5.1 contributions), ($F(1,25) = 6.01$, $p=0.02$). This, coupled with the similarity in the average size of documents between Mpathy and Npathy, and the lack of any significant difference between the topics of communication indicates that collaboration was more efficient in that participants were more able to contribute. The task set to the participants was to evaluate the system and produce a list of bugs as a group. This group activity requires discussion in order to develop the shared list of bugs, and discussion necessarily involves communication. The increased number of contributions in Mpathy is an indication that providing cues to mutual knowledge supports greater communication and so greater discussion.

More co-ordinated use of documents

Not only were there more documents created in Mpathy than Npathy, but participants also annotated more documents (Mpathy 37% documents annotated, Npathy 11% documents annotated). Moreover, anecdotally there were (non-significantly) more cross-references made between documents in Mpathy (average per group: 4) than Npathy (average: 1.2). We suggest that just by providing more cues to mutual knowledge, participants became more engaged with the artifacts in the collaborative environment. We interpret this as more focused use of the information – in commenting on a point, participants annotated the document itself rather than creating a new document and explicitly referring to the point. Again, there were less contributions in Npathy anyway, so the increased annotation in Mpathy indicates that the overall level of engagement with the material is higher when cues to mutual knowledge are provided.

Direct references to documents

Although not significant, we found that there were more direct references to documents in Mpathy (41%) than in Npathy (34%). This indicates to us that the level of mutual knowledge established during the collaboration was greater for participants using Mpathy than those using Npathy as they did not have to rely on explicitly referring to documents during discussions (which, as discussed previously, there were more of). Assumptions about mutual knowledge are key to the success and efficiency of collaboration. The key point here is that by providing 3rd order cues to mutual knowledge about who has read and annotated documents, the assumed mutual knowledge about the activity is increased i.e. there is greater understanding of what has happened in the group. Speculatively, the slight increase in Indirect references to document content in Mpathy (13%) compared to Npathy (11%) weakly supports our position that participants were relying on assumptions about mutual knowledge.

The timeline in Mpathy

Mpathy has a time based user interface, whereas Npathy's interface is based on a desktop metaphor. This distinction was introduced to allow cues to mutual knowledge to be shown in Mpathy – each participant has their own timeline on which their actions are represented relative to other participants' actions so providing a representation of the public events so far in chronological order. We argue that the improved collaboration we observed with users of Mpathy is not a product of the explicit representation of time in the interface, but rather a product of the representation of participants' activities (cues to mutual knowledge). This is because compared to Npathy's desktop interface, the timeline is extremely restrictive in the following ways which may negatively affect user performance:

- The length of the work (14 days) meant that a lot of objects (documents, memos, and annotations) were produced (average: 41.5). Users of the timeline interface would have to perform substantial scrolling to see all the documents produced as a typical window could only show about 10 documents at a time). Moreover, the cognitive load placed on users as they scroll to find documents whilst remembering where other documents is much higher than in the desktop interface of Npathy where all documents can be seen on one screen.
- The ordering of the documents in the timeline is not under user control, and it is not possible to move documents. Documents are ordered strictly by time in Mpathy's interface, whereas users may group the documents as they see fit in Npathy. It could be argued that this violates basic HCI guidelines such as supporting user control. This may also account for the slightly larger proportion of topics concerning the System functionality with Mpathy – users may have been confused to some extent when trying to impose an ordering or grouping on the documents e.g. "*just testing to see if the system permits me to change the date*". However, regardless of the usability of the timeline interface, it remains the case that Mpathy encouraged more collaboration than a conventional desktop metaphor primarily because of the third order cues to mutual knowledge embodied in the interface.

Considerations

This study focused on a very particular form of collaboration: asynchronous collaboration involving discussion and development of a single joint artifact. Whilst we believe that the notion of grounding, and the importance of providing cues to mutual knowledge is fundamental to understanding and supporting collaboration, we believe that other forms of collaboration need to be assessed in other domains e.g. synchronous negotiation activities as discussed by Clark (1996).

As discussed previously, the means of providing 3rd order cues to mutual knowledge was the timeline representation. This design allowed us to lay out all the events over time and show who had read contributions, but may have had some usability issues. In order to further strengthen our claims we need to assess other means of providing such cues as the effectiveness of such representations may vary with the nature of the collaboration. For instance, there may be novel ways to augment a more conventional desktop metaphor with indicators of who has read and accessed documents using 3 dimensional representations of the state of collaborative activity.

In the course of this study we collected a rich set of data which has much potential for further analysis. For example, we could analyze whether there is a difference in the amount of breakdown and repair that occurs in the two systems. We would expect that there would be more breakdowns in Npathy than Mpathy due to the lack of mutual knowledge about the collaboration, and so increased likelihood of misunderstandings occurring. We might also attempt to assess whether document names are used more effectively in one interface than another. Although the analysis of topics of document content did not show any significant results in this study we believe that such analyzes could yield useful results in future studies, especially if more communication channels are made available to participants.

Conclusion

This paper set out to show that asynchronous collaboration benefits from extra cues to mutual knowledge. We argued that such cues increase participants ability to contribute and promotes more focused use of information within the collaborative environment. Such findings should be of great interest to designers and developers of collaborative support systems as well as people interested in the nature of collaboration. We intend to further our research by studying asynchronous collaboration in a wider range of domains, by developing more detailed explanations of the nature of collaboration, and by iteratively informing and refining the design of collaboration support.

References

Barwise. J. (1989) The Situation in Logic, CSLI Lecture Notes Number 17, Center for the Study of Language and Information, Stanford, CA, 1989.

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.

Bryan-Kinns, N., & Healey, P. G. T. (2007). Exploring Mutual Engagement in Creative Collaborations. In *Proceedings of Creativity and Cognition*, Washington DC, USA.

Clark, H. H. (1996). *Using Language*. Cambridge University Press.

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., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In A. K. Joshi, B. Webber, & I. Sag (Eds.), Elements of discourse understanding (pp. 10-63). Cambridge: Cambridge University Press.

Dourish, P., and Bellotti, V. (1992). Awareness and Coordination in Shared Workspaces. In *Proceedings of ACM Conference on Computer-Supported Cooperative Work (CSCW'92),* pp. 107-114.

Gutwin, C., & Greenberg, S. (2002). A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work*, Vol. 11, pp. 411-446.

Healey, P. G. T., & Bryan-Kinns, N. (2000). Analysing Asynchronous Collaboration. In *Proceedings of HCI 2000*, pp. 239-254.

Hutchins, E. and Klausen, T. (1996). Distributed Cognition in an Airline Cockpit. In *Communication and Cognition at Work*, eds. D. Middleton and Y. Engeström, Cambridge: Cambridge University Press.

Issroff, K., & Scanlon, E. (2002). Using technology in Higher Education: an Activity Theory perspective. *Journal of Computer Assisted Learning, Vol. 18*, pp. 77-83.

Johnson, H. & Hyde, J. (2003). Towards Modeling Individual and Collaborative Construction of Jigsaws Using Task Knowledge Structures (TKS). *ACM Transactions on CHI, Vol. 10, No. 4,* December 2003, pp 339–387.

Kindberg, T., Coulouris, G., Dollimore, J., & Heikkinen, J. (1996). Sharing objects over the Internet: the Mushroom approach. In Proceedings of GLOBECOM '96, pp. 67-71.

Kindberg, T., Bryan-Kinns, N., and Makwana, R. (1999) "Supporting the Shared Care of Diabetic Patients." Proceedings of ACM GROUP '99, Phoenix Arizona.

Koschmann, T. & LeBaron, C. D. (2003). Reconsidering Common Ground: Examining Clark's Contribution Theory in the OR. In *Proceedings of ECSCW 2003*, Helsinki, Finland

Olson, J. S., Olson, G. M., Storrøsten, M., & Carter, M. (1993). GroupWork close up: A comparison of the Group Design Process With and Without a Simple Group Editor. *ACM Transactions cm Information Systems, Vol. 11, No 4*, pp. 321-348.

Robertson, T. (1997). Cooperative Work and Lived Cognition: A Taxonomy of Embodied Actions. In *Proceedings of ECSCW*, pp. 205-220.

Watts, L. A. and Monk, A. F. (1998). Reasoning about tasks, activities and technology to support collaboration. *Ergonomics, Vol. 41, No. 11*, pp. 1583-1606.

Acknowledgments

This work is supported by EPSRC grant Engaging Collaborations GR/S81414/01, and previous work on the Mushroom project funded by EPSRC grants GR/L14602 (1996-97) and GR/L64300 (1998-2000).