

Informed Opportunism as Strategy: Supporting Coordination in Distributed Collaborative Writing

Eevi E. Beck [†] *

[†]University of Sussex, UK

Victoria M. E. Bellotti *

*Rank Xerox Cambridge EuroPARC, UK

Abstract: There is little understanding of how distributed writing groups manage their collaboration and what kinds of support are most useful. The paper presents three case studies of distributed collaborative writing groups in academia. The process evolves over time, constantly adapting to changing circumstances. Co-authors offer and make use of a range of information. Their subsequent opportunistic use of this information to make appropriate ad hoc decisions in new circumstances, appears to be essential to achieve flexibility and coordination. We call this *informed opportunism*. We identify design implications for support tools for distributed collaborative writing.

1 Introduction

The rapid development of communications technology is making it easier and more cost effective for many people to communicate with others over a wider area than ever before. Some find it convenient or necessary to work closely with colleagues who are great distances away, and to collaborate on documents with them. The chances are increasingly that co-authors of documents are distributed over a wide area and must use that same telecommunications technology to support communication and management of writing activity.

Research on co-authoring is increasing in the humanities (Lunsford and Ede 1990; Forman 1992) as well as in CSCW. An increasing number of interview studies (e.g. Posner 1991; Rimmershaw 1992) and surveys (e.g. Lunsford and Ede 1990;

Couture and Rymer 1991; Chandler 1992; Beck 1993) are becoming available. Few longitudinal studies of the process of collaborative writing exist (although see Law and Williams 1982; Riley 1983; Couture and Rymer 1991; and Plowman 1992; and a growing number of self-reports of the experience of co-authoring). Co-authoring can be defined in many ways (Lunsford and Ede 1990), giving different perspectives on what the practices of co-authoring are (Couture and Rymer 1991). The assigning of authorship to academic papers is in itself a subject of study (*e.g.* Trimbur and Braun 1992).

We are interested in 'close' co-authoring, as opposed to, for example, editing, and define this as the writing of documents by two or more persons, where the names appearing in the author list are those of people who considered themselves involved in the writing of that document. We feel that for distributed co-authoring groups, *i.e.* those working together across geographic distance, there is a particular potential for appropriately designed computer-based tools being helpful.

Over the past few years there has been a rapid growth in the number of collaborative writing tools, *e.g.* QUILT (Leland, Fish, and Kraut 1988), GROVE (Ellis, Gibbs, and Rein 1990), SHREDIT (McGuffin and Olson 1992), and PREP (Neuwirth, Chandhok, Kaufer, Erion, Morris, and Miller 1992); and hypermedia co-authoring tools, *e.g.* COAUTHOR (Hahn, Jarke, Eherer, and Kreplin 1991). Some are designed to support writing together both at the same time and at different times, *e.g.* SASSE (Baecker, Nastos, Posner, and Mawby 1993). Few are designed for distributed co-authoring, an exception being MESSIE (Sasse, Chuang, and Handley 1993).

We have studied co-authoring in academia and have, however, found no evidence of use of special collaborative writing tools, except among those close to the development teams. Instead, face-to-face meetings, conventional workstation and printing technology, telephone, facsimile (fax), and post appear to be the means by which information is exchanged and activities managed. This is consistent with a general concern, *e.g.* with Kling (1991), that the use of CSCW tools has not proliferated at a rate corresponding to the increasingly sophisticated range of systems available. This problem has been attributed to a poor understanding of the way in which groups collaborate (Kraemer and Pinsonneault 1990; Grudin 1991).

An early indication that a distinction between discussion and annotation may not be useful for all co-authors (see Mhashi, Rada, Beck, Zeb, and Michailidis 1992), prompted our interest in whether other distinctions proposed in the analysis of writing make sense to users of the systems¹, or whether inappropriate assumptions are finding their way into collaborative writing systems. Does co-authoring proceed through 'stages' which it is meaningful to design to? COAUTHOR is one of several systems which has a strong separation of idea generation from text writing. Here, despite "no rigid phase model [being] intended" (p. 81), strong assumptions of what

1 Our concern here is not with whether such distinctions are appropriate in the analysis of writing or collaborative writing, nor with whether they provide useful frameworks for thinking about writing support, but, rather, with whether wholesale application of analytical and theoretical distinctions into support tools, distinctions such as those based on Flower and Hayes' influential cognitive model of writing (*e.g.* Flower and Hayes 1980); encourage a task oriented approach which may prove unnecessarily restrictive to users. Note that Flower herself in a later paper calls for writing research now to "explain how context cues cognition" (Flower 1989, p 1)

the phases are, and the order in which work should be done, are evident: "The first phase of *idea processing* determines the issues (...). During the second phase of *document design* a formal document structure has to be set up and associated with the conceptual items from the idea processing phase." (Hahn *et al.* 1991, pp. 80-81, orig. italics). Explicit role support has been called for *e.g.* in Baecker *et al.* 1993. But do co-authors allocate responsibilities *pre hoc* and stick with those decisions so that choosing and enforcing 'roles' is useful? The co-author roles arrived at in Posner 1991 from interviews with co-authors, were derived from *post hoc* analysis of what activities the co-authors had spent most of their time on and were not necessarily what they would have chosen initially (indeed, the final role had changed from initial expectations in one of the cases), whereas explicit role support presumably would require co-authors to commit to roles early. In what sense do writing groups devise 'strategies'? Sharples 1992 gives as pre-existing 'strategies for distributed collaborative writing' patterns of task distribution ('parallel', 'sequential', 'reciprocal') which have little empirical basis. Our concern is that unhelpful assumptions about co-authoring may become reified and find their way into system design, repeating the problem identified by Tatar, Foster and Bobrow (1991): in examining reasons for failures of the COGNITER system, Tatar *et al.* found that incorrect assumptions about human communication had become enshrined in the design. They concluded that "[i]n a field that is as new and as complex as computer-supported cooperative work, [...] highly directed studies need to be augmented by other approaches such as undirected observation" (Tatar *et al.* 1991, p. 207).

The aim of this paper is to document practices in distributed co-authoring groups which we have observed in academia which do not support the above assumptions: salient features are not the carrying out of tidy agreements, but great flexibility and context sensitivity with which co-authors interpret information and situations and come to decisions about appropriate courses of action, even to the extent of unilaterally **contradicting** agreements. We call this strategy *Informed Opportunism*².

Below, we summarise the findings of two preliminary studies which looked at organisational aspects of co-authoring. These studies provided the impetus for the subsequent longitudinal case-studies. Then we present the case-studies themselves, examining in more detail the dynamic changes in the co-authoring process, with a particular emphasis on two basic questions about the nature of and support requirements for the coordination of individual activity within distributed co-authoring groups. Finally, we present design recommendations based on our findings.

2 Preliminary studies

Two studies of group organisation and collaboration aspects of co-authoring practices were undertaken. Initially, semi-structured interviews of 1-2 hours length were conducted with about ten co-authors. Issues which appeared to be the most impor-

2. 'Opportunism' is intended in the sense of "being guided by what seems possible, or by circumstances, in determining policy; preferring what can be done to what should be done" (Oxford Advanced Learner's Dictionary of Current English, 1974 edition) We refer to this as a 'strategy' only in the sense of "skill in managing any affair" (Oxford Adv. L.'s Dict. of Curr. Eng., *ibid*), not *e.g.* the systematic making and execution of plans.

tant to the interviewees were reshaped into an exploratory questionnaire survey which was distributed to 23 academic researchers.

The respondents answered 24 multiple choice and rating questions, as well as supplementary open ended questions, with respect to one instance of co-authoring they had experienced. The questions addressed aspects of co-authoring practices ranging from distribution of responsibilities and discussions of the organization of the group, to membership changes, relative contributions by co-authors, motivation, perceptions of success and notions of collaborative writing.

Examples of questions asked are: Except for right at the beginning, were you at any point ever unsure about who were going to be the co-authors? (Replies allowed: Yes/No/Don't know; elaboration requested for yes-answers). What is, to you personally, the purpose of taking part in collaborative writing? (open-ended question). The main findings of the interview studies and survey are summarised in the following sections (for further details on the survey, see Beck 1993).

2.1 Group membership and roles

One third of the survey respondents indicated that membership of their writing group had changed while writing. Five respondents (23%) reported being unsure at some point who the co-authors were going to be. Thus group membership in the general case cannot be taken to be static, but is in many groups dynamically changing. Group leadership was attributed in equal proportions to a self-appointed leader or facilitator; an agreed leader/facilitator/project manager, and 'no-one'. This suggests that a system designer would be unjustified in assuming the presence of an agreed 'leader' in an academic collaborative writing group. The extent to which responsibilities were clearly divided between co-authors appeared to vary, with some evidence that it was more common for there to be overlapping areas of responsibility than not. Co-authors in the interview study reported frequent departures from agreed areas of responsibility in response to the progress of the joint writing.

2.2 Plans and goals

All survey respondents indicated that a very important measure of success was that the resulting document was acceptable to themselves. There was less agreement on quality, good communications, and getting on well with each other being important, although these also scored high means. The contents and structure of the documents were not reported planned in advance so much as discussed during the writing. The question of the importance of 'Close adherence to an initial plan' received extremely varied ratings, with a medium to low mean.

2.3 Implications

The preliminary studies indicated that academic co-authoring groups may be subject to a complex mix of environmental influences, to the extent that roles, goals, and strategies must be regarded as subject to unpredictable change at any time. One

interesting issue is, if there is such unpredictability, how do co-authors achieve the necessary coordination of activities to produce a document together? This is what our case studies set out to explore. We shall see how this suggests approaches to answering the question of appropriate system support. For example, the observation that the nature of the group and its mode of operation evolves dynamically in response to a multitude of factors relating to both group and individual member's contexts, gives an indication that a tool which is to fit in with this must incorporate great flexibility to accommodate the adapting process.

2.4 Communication and coordination: the research questions

A crucial point, we feel, is the manner in which individual work is related to the group as a whole. Co-authors make autonomous decisions when working alone, under changing and unpredictable conditions, which the group cannot foresee or plan for. However, in close collaboration each author must maintain awareness of and offer information about the current state of affairs to others, so that work is coordinated and appropriate contributions are made.

Communication mechanisms would seem critical to coordination of the collaborative writing endeavour. For distributed authors it is these mechanisms which are most likely to be particularly problematic. We would like to understand more about:

- How a document comes to exist and evolve into its final version.
- How co-authors manage access to the document.
- How co-authors decide what and when to write.

In each case we focus on information exchange and ask what co-authors inform each other of, and how.

In the following section we describe our investigation of these issues. Section 4 addresses the design implications.

3 Three case studies

Three observational case studies were conducted of *in vivo* distributed collaborative writing, *i.e.* as it was taking place in everyday work. The co-authors were employed in research in academia and industry. The completed documents were to be papers presenting research findings for publication in academic journals or conferences. The length of the joint writing ranged from three weeks (plus a post-review, final camera-ready copy preparation effort) to over a year.

3.1 The groups

Group A

Context, document and members: There were three co-authors, here referred to as A1, A2, and A3. Co-authors A2 and A3 were in the same building, whereas A1 was at a different institution 2-3 hours' drive away. The three formed the core of a research group which had already published one journal article. The document fol-

lowed in this case study was another journal paper reporting on further research findings. There was no deadline. A1 and A2 were working on two other papers as well as a book concurrently with this paper.

Communication and supporting technology: A3 had been anticipated fully involved, but rarely took part in meetings, even at his own institution. Instead, A1 and A2 produced most of the paper without him. A1 and A2 were in regular contact, travelling to meet face-to-face 1-2 times per month, and telephoning 1-2 times per week. They made use of stand-alone text editors, diskettes, and occasionally fax transmission (3-4 times over four months). Other technologies available included electronic mail (email), which was viewed as potentially useful but never used.

Group B

Context, document and members: There were two co-authors, B1 and B2, separated by great distance and an 8 hour time difference. They had done research together, partly while in the same place, but mainly at a distance. The paper was to describe the results of their joint work and was to be submitted, in camera-ready format, to a conference. During the three weeks while most of the writing was done, there was considerable time pressure to meet the main deadline for dispatch (after the paper was accepted, further modifications were made).

Communication and supporting technology: B1 and B2 met face-to-face very early on, before the writing really started, but not while writing. There were daily email exchanges. Telephone calls were 1-2 per week during the busiest 3 weeks. Fax was used once, to send publishers' formatting instructions. Both authors used a sophisticated single-user text editing package. They transferred the entire document between them up to three times per 24-hr period, using electronic file exchange on networked workstations. This was by far the most intensive of the three groups.

Group C

Context, document and members: Two co-authors, C1 and C2, worked in the same department while developing the first draft, but were then separated by significant geographic distance (5 hours time difference). There was no deadline.

Communication and supporting technology: Prior to separating, printouts of the entire document were transferred twice. At separation, C2 took a copy of the file on a diskette. After separation, email was the only means of communication until a visit by C2 to C1 (for purposes other than this paper). The cost of telephoning was considered prohibitive. The writing is ongoing, and only the group's initial work is considered in this paper.

3.2 Method

In each case-study, interviews were conducted with the co-authors as soon as possible after their decision to write a paper, and after the completion of the writing. There were frequent follow-up interviews, mainly over the telephone. Interviews were audio recorded, and conducted with co-authors on their own. In addition, some




Figure 1. Growth in group B's document size over time. The horizontal axis represents time in days at which a version of the document was exchanged from co-author B1 to co-author B2 and vice versa. Dashed lines indicate a compressed timescale. The vertical axis indicates size of the document in blocks (one block is approximately 1000 characters' worth of data). Each version is represented by two bar heights to indicate the size of the text-only document file and the formatted document file. (Note that the formatted sizes exclude a large figure in the original versions (about 30 000 blocks) which the coauthors inserted into version 2, removed again (because it was slowing down their system) for version 6, and reinstated for version 18.) Two gaps in time are indicated by broken lines: between versions 1 and 2 (one week), and between versions 21 and 22 (approximately 4 months).

observation of co-authors working together was possible. Email messages, as well as notes on and intermediate versions of the document, were stored. In group B, all communication between the co-authors about the paper was recorded.

In the following three sections we report on our findings on what and how distribution of information took place. We then summarise our findings and address the question of desirable features of support tools for distributed co-authoring.

3.3 Document creation and evolution into its final version

Creation

In both groups A and B, an explicit decision to start writing a paper existed before writing started. Group B met to discuss the topic and content of their document before they began writing it. In group C, on the other hand, there was a gradual process by which an initial note was developed into a long enough document that the authors decided to turn it into a full-length publication. For all three case-studies the shared document coming into existence was a gradual realisation; it was not an event which even in retrospect could be clearly determined.

Evolution

For group B, we were able to store a complete set of versions of the document exchanged between the co-authors. Figure 1 shows the development of the paper in terms of the sizes of each of the 24 versions they exchanged.

The bulk of the writing took place over 22 days. The gaps between versions 2 and 3, and 18 and 22, were due to absence of B1 or B2; in each case, the other one

went on writing. So, although the last exchange before the deadline was day 17 when B1 went away on a trip, B2 continued writing up to day 19. On day 11 the network connection was not working and thus no exchanges were possible.

A high level of activity was evident until B1's departure, with a peak at three exchanges of drafts on day 16. The deadline for sending off the document was day 19, when co-author B2 submitted version 21. (In Figure 1, version 21 is shown on day 22, when B1 returned and saw a copy). After notification of paper acceptance, versions 22-24 were produced as a final revision before publication. These were exchanged over two days approximately four months after initial submission.

Figure 1 shows that in group B, the document reached its full size early: already version 7 was of a similar size to the last versions. Examination of the actual version changes revealed that much of the total effort went into rearranging and rephrasing existing text, as opposed to generating new text.

Concern about formatting emerged early in B1's and B2's communications. They put in headers, some with mock section contents, which helped organize existing text and apparently acted as placeholders and reminders of actions to be completed in the future. Headers were also important aides in communicating summaries of changes, where they were extensively used as coordination points or shorthand references to the section contents: "[The new version] *has the* [required] *format (hopefully), the figure moved to the last page (since it looks like it will require manual placement), a rewrite of the* [section name] *stuff, and various other small modifications.*" (email message accompanying version 6).

Some formatting changes also show up in the graph: between versions 5 and 6, the document was re-formatted from single-column to double-column, resulting in an increase in the formatted size but not the text size.

In producing group B's document, no significant restructuring was done from version 13 onwards. Interestingly, most other changes—additions, deletions, and rewordings—occurred throughout. The volume of changes trailed off after version 18.

For group A, the process of writing was closely intertwined with other work including writing of other publications. Their way of working together was diverse, changing throughout the process. At face-to-face meetings, they would often sit down by the word processor and compose or review parts of the document together. Twice, an electronic file was exchanged by storing it on a floppy disk and handing this over at face-to-face meetings. Between meetings, they would normally work on separate sections, which one of them (usually A2, who was more experienced with the word processor) would incorporate into the current version. During the three months of writing, the document evolved largely through parts being passed between A1 and A2 (rarely the whole document), typically every week or two.

The rate of communication in group C was the lowest of the three case-studies. Over one year the document was expanded from an initial 3-4 page version written by C2 without a view to producing a full-length paper, to an expansion by C1 into 20-odd pages; in retrospect referred to as the first draft. C2 had done related writing for other purposes and incorporated this into a second draft, which was discussed

when they met six months later. After one meeting discussing the paper, C2 made amendments and left the paper with C1 to review, which C2 did after some months. Currently the paper is back with C2, for what was hoped would be a final revision.

3.4 Managing Access to the Document

Changing membership and contributions

Group membership means access to the document. In these case-studies, there was evidence of tacit expectations existing of levels of contribution from co-authors (all groups), and of dissatisfaction (group B) or slowdown in progress (groups A and C) when expectations were not met. In group B, the expectation that both co-authors work enough was an issue that kept being referred to in email messages, e.g. “*you would’ve noticed that I’ve actually done some work*” (in reference to an email message that went missing).

We did not see evidence of specific roles being allocated. However, in group A, the two main co-authors initially split the sections fairly strictly on expertise. But gradually this specialization broke down; each co-author, instead, opportunistically doing what they could, often taking into account how busy the other co-author was.

Timing and organisation of document access

The available technology and the respective work rhythms (such as determined by time zone differences or other work commitments) seem to be important factors in choosing access strategies, as well as time available, or urgency of the writing. In groups B and C, explicit agreements on how to access the document were followed most of the time, whereas in group A, no such general agreement on document access was observed.

Opportunistic decisions on document control were most evident in group A. For example, one version of a document was exchanged between co-authors by handing over a diskette when the co-authors were meeting for other reasons (a seminar).

In group B, an alternative document control strategy of parallel work and later merging emerged to allow both co-authors to work despite a technical communications problem. Furthermore, B2 once changed the document despite B1 “having” control. B2 knew B1 was tied up with other work, and that there was no system enforcement of the control. This opportunistic breach of agreement caused no problem to B1 when told; he knew B2 had had the knowledge to make a good decision.

In group C, the document was initially passed between the two co-authors in its entirety, such that each had complete control. This strategy was later changed to a plan for C2 to be in charge of combining sections, as indicated above, and still later readjusted in response to slower progress than anticipated.

3.5 Deciding What and When to Write

Co-authors may agree on writing strategies, but these are not necessarily adhered to. None of the groups had an evident leader; instead individuals made agreements as

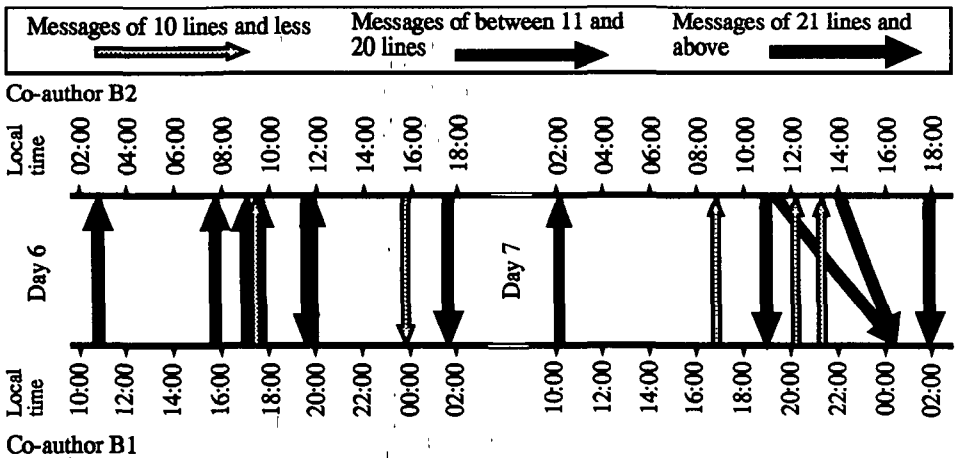


Figure 2. A sample of the email exchanges between the two co-authors in group B. Arrows indicate directions and relative sizes of email messages. (Sizes are given as numbers of text lines in the message, excluding header lines).

required, and made autonomous decisions and took opportunistic action in response to unforeseen developments. Agreements on who does what, when, were made and broken in all three groups, without necessarily causing problems.

Communicating editing and organisational information

In group B, the process was characterised by periods of intense communication when possible (using email, file exchanges and telephone calls), interspersed with longer periods of individual, autonomous work when one co-author was away.

Figure 2 shows an example of group B's email exchanges during days 6 and 7, a time when the document size was rapidly expanding, and with the highest rate of messages. Sending and reception times are indicated in the local time of each co-author. It should be clear that transmission time was usually almost instantaneous. Two messages from B2 on day 7 were held up, and crossed with other messages, which went undetected for a while, but in this case caused no problem. Once, however, a missing message did cause problems and subsequently some double-checking was done on the 'phone of whether email had got through. This suggests that the unreliability of email may cause problems for even experienced email users.

It is interesting to note, in Figure 2, that each co-author would send several messages apparently without waiting for a response. This was a typical pattern of exchanges in group B. Specific coordination information not requiring a reply was often contained in short messages: *e.g.* that a new version of the document was ready for handing over. This information was important, as it marked transfer of responsibility for the document—of who "had" it. Note also that the email messages are relatively short—two-thirds of the messages have less than 20 lines of text. This was fairly typical for the email exchanges in group B³.

In group A, towards the end of writing, A1 and A2 had agreed that A1 should look over the near-final version and provide A2 with final changes. But after waiting

3. This contradicts a claim in Sharples 1992 that long communication is required for intense collaboration

for a while without any notification of changes, and being aware that A1 was particularly busy in that period, A2 opportunistically made the last changes himself, in order to get the paper submitted sooner. It is interesting to note that both A2, and B2, in the similar situation above, were able to do this because access restrictions were not technically enforced; only socially. A1 or B1 were only notified after the event, but no disquiet about this breach of agreed control strategy was evident.

Changes were not always communicated and this sometimes lead to problems. Communication failures could result in co-authors becoming disgruntled with one another, and could mean that precious effort went to waste. On one occasion an important email message from B2 failed to get to B1. It set out a proposed structure and some contents for the paper and was sent off just before a week's absence by B2. Despite regular communication in the intervening period, it took a few weeks to discover the loss, by which time the message and the effort that went into it was largely wasted. On discovery of the loss of the message, B2 was clearly disappointed: "*(Jeez), B1, and I worked so hard to get this out before I left*".

Group C had a different problem, in that an expectation of no feedback in case of progress meant that at one point, C1, having heard nothing, was half waiting for a second draft to turn up at any time when, in fact, there had been no progress.

Making changes visible

In group B, co-authors made their own changes visible to their co-author in a number of ways: telling them explicitly in summaries; putting in devices to attract attention in the text itself (comments in italics and square brackets), and by the removal of such devices.

Both co-authors would, at times, make comments in the text of the document itself. There might be replies to such statements or queries, and a dialogue would take place through the document versions. The ability to attach a comment to the point in the text it referred to, allowed efficient use of deictic references. This was extensively used, *e.g.* in version 14, in mid-paragraph: "... is a common activity. *[B2: Why? I don't think this is intuitively obvious. There are implications in the next paragraph but maybe something would be useful here?]*" (orig. brackets and italics). Ideas could be suggested: "Based on *[B1: something or other]*, we have ..." (in versions 12 and 13; orig. brackets and italics).

Comments were status indicators; a sign that further work was needed. Comments would remain until a solution was implemented. Thus *e.g.* the second comment above remained for two versions until replaced by "Based on [our experiences], we have ..." (version 14, our brackets). Another way of making changes visible was employed when B2 had made changes to sections of the document while B1 "had control" over the document (see above). B1, having made other changes to the same sections, left the old version of the sections in question at the end of the document for B2 to see, as a historical trace of where the new version had come from. (This shows up in Figure 1 as version 12 being significantly longer than 11. Note that this was deleted for version 13 by B2.) Finally, both co-authors would, in email messages, regularly summarise the changes they had made to the document.

We thus see how these co-authors achieved compact communication of changes by employing both a textual and historical context to carry part of the meaning.

Balancing work on the document with other activities

The co-authors were continuously having to balance work on the joint document with other activities. In group B, where the effort to produce the paper was most concentrated, there was an expectation of very high prioritising of work on the paper. In group A, the writing took place over a longer period and more other work was done in parallel. In group C, the proportion of time spent on activities unrelated to the collaborative writing was very high for both co-authors.

Co-authors made efforts to find out about and inform others about competing activities, and continuously made adjustments to expectations of the other's work. The exchange of detailed information about each others' program was particularly evident in group B. But also in group A, co-authors A1 and A2 knew a lot about each other's program (*i.e.* they had told/asked each other). *E.g.* A1 knew that A2 would be having a small operation one week and did not expect him to make any changes to the paper. In group C, C1 carried on with other work while remaining ready to prioritise the joint work, should C2 send him a new version of the document, showing great flexibility in adjusting to C2's pace of work. However, lacking the information that C2 was busy with other work, caused C1 to be unsure whether to set time aside from other work for this paper. Note that C1 initiated communication with C2 at a point when other work was causing less pressure on him.

4. Summary of findings and design recommendations based on case studies

In this section we summarise our case-study findings and address the design implications of the work reported above. Our data does not allow us to generalize, but does provide some pointers to issues to address. In particular, our observations of great variety in writing practices, a finding consistent with those of Rimmershaw in her interview studies (Rimmershaw 1992), we feel poses a particular requirement for designing for **flexibility**. Agreed protocols may be adapted, or individual authors may make autonomous decisions as to whether it is appropriate to break agreements. This is not done in an irresponsible fashion; rather, it is in response to continuously changing circumstances, unforeseen events, technological breakdowns, and so on. Awareness by the designer that groups may change over their lifetime, might mean, for example, allowing the way that document versions are managed, any access restrictions are enforced, and how annotations are done, to be dynamically changed by users. Designing for this kind of flexibility will, we believe, prove crucial to highly usable co-authoring tools.

The design recommendations presented below are an initial attempt at addressing what such flexibility might mean in terms of design issues. These are primarily intended to pose questions, not answers, about academic and similar co-authoring.

4.1 Coordination of editing activity

Integration with standard platforms

Shared documents may evolve from private documents, files from electronic communications and so on. Thus, support tools which allow existing text in other forms to be easily incorporated, may be useful. (Many systems currently do this.)

Formatting

Co-authors may want to make formatting changes throughout. Easy-to-change formatting can be used in a flexible way not only to help visualize the final document, but can also support a range of signals among co-authors about interpretive context, e.g. by making parts which are **not** intended to be in the final version stand out. Computer-based tools could support flexibility by allowing exchange of the document in a formatted form throughout the writing process.

Grounding communication in the document

It appears that for some changes it is important to communicate their rationale. Tools might usefully support linkage of communication about the document to changes in the document. Such linkage should foster common grounding of communication between authors, improving the comprehensibility of discussion about the document and allowing efficient use of deictic references.

Communicating changes

The co-authors regularly drew others' attention to some of their changes by commenting on the location, and, frequently, their nature. Co-authors also communicated an intended incompleteness of what had been written through various cues (such as bracketed comments). This may be important in indicating a willingness to have changes made. Support tools could reduce overheads of highlighting selected changes by providing easy means of referring to change locations and optionally allowing notes about the changes to be attached. PREP's "flexible diff" (*ibid.*) highlights changes of specified sizes (sentence, word, paragraph).

Author information

Decisions about when and what changes to make to a section were at times influenced by who had written it. A co-authoring tool might provide lightweight support for knowing who did what, which a user group may turn on and off. Several systems, including SASSE (*ibid.*), provide author information automatically.

History information

Editing may take place over a brief or extended period. Tools could incorporate lightweight support for making judgements about the completeness of, and confidence in, parts of the document. For example, an editing history may be presented as a quick replay of changes. This might cue co-authors about the context of particular changes, and might conceivably reduce the problem of information overload.

4.2 Coordination of responsibility

Planning and status information

Authors must decide which part of the document they can most usefully work on and what to contribute in order to best fulfil goals and avoid duplications or omissions. No individual 'leader' can be assumed to take the lead in this. Also, plans are not necessarily carried out even if explicitly agreed. Tools may be of more benefit to co-authors by supporting general exchange of information about future expectations; about goals achieved or attempted, *etc.*, than by requiring users to follow plans previously made. For a structuring and planning tool, such as GROVE (*ibid.*), particular attention may need to be paid to making previous work easy to change.

Information about ongoing activity and progress

Co-authors may be engaged in a range of other unrelated activities which make demands on their time and attention. To organize their own activities and set time aside for making contributions, they are dependent upon knowing about each others' activities. Support tools could help foster group awareness of progress, *e.g.* by broadcasting automatic notifications as to when particular authors begin editing activity, and perhaps about what part of the document is being edited. Such automatic information must be possible to turn on and off easily. SASSE (*ibid.*) provides a document overview tool which indicates who is working on the document and where.

Control of access

At times, for maintaining consistency or ensuring that the best qualified person carries out certain tasks, it can be useful for one author to have exclusive access to part or all of the document. This may be successfully socially enforced, but support tools may feature facilities to alert co-authors to the status of document parts with respect to currently agreed access rights. MESSIE (*ibid.*) permits one user at a time to have editing access to the document and sends an email message with information about who is editing it to anyone else attempting to change it. This, however, removes from the group the possibility of working on parallel versions and merging them later, and may discourage short or intermittent working.

Roles vs. responsibilities: Inform, rather than constrain

Social control can be, it would seem, highly effectively used to manage the rights and responsibilities of each participant in a writing project. It allows areas of responsibility to be in continuous change without overheads, and avoids the problem that desired access restrictions may not coincide with predefined roles (Dourish and Bellotti, 1992). We therefore believe that rather than providing generic co-author roles for tools to support, designers might usefully seek to support co-authors' exchange of information to help them make their own judgements about appropriate contributions. Some such information may be appropriate to collect and make available automatically, other not; more research is needed on this.