

Information and Context: Lessons from a Study of Two Shared Information Systems

Paul Dourish, Victoria Bellotti, Wendy Mackay and Chao-Ying Ma

Rank Xerox Cambridge EuroPARC
61 Regent Street
Cambridge UK
CB2 1AB
+44 223 341500
dourish@europarc.xerox.com

ABSTRACT

With the increasing ease and power of computer networking technologies, many organisations are taking information which was previously managed and distributed on paper and making it available electronically. Such shared information systems are the basis of much organisational collaboration, and electronic distribution holds great promise. However, a primary focus of such systems is on the ease of information *retrieval*. We believe that an equally important component is the problem of information *interpretation*, and that this interpretation is guided by a *context* which many electronic systems do not fully acknowledge.

We report on a study of two systems, one paper-based and one electronic, managing similar information within the same organisation. We describe the ways in which information retrieved from these systems is interpreted subjectively by individuals, and point to some of the factors contributing to this interpretation. These factors, together making up the context of the information, are of critical importance in the design of successful electronic shared information systems.

KEYWORDS

Collaborative systems, cooperative work, information systems, information context, coadaptive systems.

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

COOCS'93 - 11/93/CA, USA
© 1993 ACM 0-89791-627-1/93/0010...\$1.50

INTRODUCTION

Information systems, electronic and non-electronic, are ubiquitous in organisations. Databases, diaries, calendars, phone lists and organisational records facilitate communication and collaboration within working groups. With the advent of widespread computer networking, more of these information bases, traditionally distributed on paper, are being made available to groups electronically.

However, the transition from traditional to electronic media for the management of this information is not straightforward. The adage that “an information system is only as good as the information it holds” is only partly true. This is because, in addition to being merely retrieved, the information must also be *interpreted* so that an individual can decide how to use it. We are interested, here, in looking at this process of interpretation—at how and when it takes place, at the resources which support it, and at the implications it holds for the design, deployment and evaluation of electronic shared information systems.

An Example

A simple example will serve to illustrate how we use contextual interpretation in everyday information management activities. Consider the difference between a telephone number written on a scrap of paper, one published in a public telephone directory and one printed on a business card. As we deal with the same information in these various forms, we will process it, judge it and understand it differently. The difference is not in terms of the underlying data, but in terms of the context of its delivery and, in this case, primarily the formality of the medium. Similarly, the *source* of the information (from the individual concerned, from a colleague, and so on) may also affect the way the information is interpreted.

Our focus has been on the *context* which information carries with it within a system, and which can be employed by the users as part of the evaluative and interpretive process. We have primarily been focussing on a particular case study within our organisation. EuroPARC has two parallel shared information systems, the Calendar and Khronika [8]. These two systems both provide electronic and non-electronic access to largely similar information bases, and they are both available to all members of the laboratory. We have examined the ways in which these two systems are used, and report here on some of the differences in usage patterns which reveal aspects of the information context, and have significant design implications.

RESEARCH STUDY

Two Information Systems: The Calendar and Khronika

The Calendar: From its earliest days, a regular feature of life at EuroPARC has been the distribution of "The Calendar". This document lists details of the whereabouts of laboratory members over the coming weeks, as well as information about visitors, upcoming meetings and seminars, and events such as holidays. The Calendar is distributed weekly both by electronic mail and in paper form, and is delivered to every lab member as well as interested parties outside of the lab.

Khronika: EuroPARC also has an electronic event server which allows users to browse or update a database of events from any workstation. Khronika extends the basic idea of an "event database" mechanism in two ways:

1. The notion of "event" is generalised to include unscheduled, electronic events as well as scheduled calendar-style events. So, as well as seminars and meetings, the arrival of electronic mail or actions within the RAVE media space [4] can be processed as events.
2. Khronika users can create *daemons* (active agents) that seek out information of interest to a user and generate "reminder" events. Reminders may send electronic mail about upcoming events, spawn other computational processes, or generate synthesised speech or non-speech audio cues [2] through the media space.

It should be emphasised that Khronika is a shared system for calendar-style (and other) events; it provides no facilities for automatic event scheduling (*e.g.* [1], [6]). Khronika deals merely with shared *information*; there is no notion of shared *action* (such as automatic scheduling of meetings). Although its reminder mechanism is active, the emphasis on information makes comparisons with the paper Calendar valid.

Khronika and the Calendar parallel each other in interesting ways. Clearly, they serve some of the same purposes, and the

information which they collect overlaps to a considerable extent. In fact, Khronika was originally envisaged as an electronic form of the Calendar. However, a more detailed examination of the ways in which they collect and distribute information reveals important differences. These differences, and their causes, are the principle focus of this paper.

Data Collection

We studied the use of these systems through automatically-collected data, interviews and a questionnaire.

Context: The two systems which we are investigating are in use at EuroPARC, a Rank Xerox research lab in Cambridge, UK. EuroPARC is a small lab (around 30 staff), on three floors of a single building. While there are many face-to-face encounters, electronic communication is also very common, including extensive use of electronic mail, and an audio/video "media space" communication infrastructure.

Participants: The participants, EuroPARC's staff, include research, management and administrative staff with differing requirements for scheduling and coordination. The research staff are drawn from a variety of disciplines; they are divided approximately into thirds with backgrounds in computer science, psychology and social science, so computer skills vary widely. Most have extensive experience with the Calendar, which has been in use for over four years, and Khronika, which has been in use for over two years.

Data: 20 members of the lab responded to a ten-page questionnaire which asked about technical background, job requirements for scheduling and coordination, ways in which the Calendar and Khronika are used and an open-ended section for views on each system. Automatically collected data included snapshots of the events currently in the database and statistics on usage. We also conducted informal interviews with a number of individuals, including the administrator responsible for the Calendar.

In order to understand the systems as they exist in practice, we will begin by looking at them individually. We will first examine how the Calendar functions in the context of people's everyday activities, and then turn our attention to Khronika. We will then compare the practices surrounding these systems, looking at specific differences in usage. In particular, we shall focus on the ways in which people make explicit distinctions and choices between the two systems, and the reasons for making such choices. These distinctions and decisions are the basis from which we can begin to discuss the contextual information which the two systems embody and the ways in which it is employed by the user communities. From this, then, we will go on to discuss the implications of this contextual interpretation for the design of shared information systems generally.

EuroPARC Calendar Week of 10 February 1993

Volume: The PARC U 22930 1993

HERE AND AWAY THIS WEEK

Not in table: computer workshops, Laser and Robotics work shop

Activity	M	T	W	T	F	Sat	Sun	Notes
Admission	-	-	-	-	-	-	-	PARC, Open University, University of London
Ballroom	M	T	W	T	F	-	-	PARC
CAFÉ	M	T	W	T	F	-	-	Cambridge University, Computer Lab
Conferences	M	T	W	T	F	-	-	PARC
Course	M	T	W	T	F	-	-	Cambridge University, Computer Lab
Demos	M	T	W	T	F	-	-	PARC
Discussions	M	T	W	T	F	-	-	PARC
Physic	M	T	W	T	F	-	-	PARC
Course	M	T	W	T	F	-	-	PARC
Harper	M	T	W	T	F	-	-	Workshop at ITC, Watlington, Oxford, City, IT, Mathematics
Health	M	T	W	T	F	-	-	PARC
Workshop	M	T	W	T	F	-	-	Open University
King	M	T	W	T	F	-	-	PARC
Lecturing	M	T	W	T	F	-	-	PARC
Longhouse	M	T	W	T	F	-	-	PARC
Loft	M	T	W	T	F	-	-	PARC
Machinery	M	T	W	T	F	-	-	University of Surrey
Materials	M	T	W	T	F	-	-	IT Mathematics
Me	M	T	W	T	F	-	-	PARC
Midway	M	T	W	T	F	-	-	IT Mathematics on Thurs, Visions
Workshop	M	T	W	T	F	-	-	Workshop at ITC, Watlington
Others	M	T	W	T	F	-	-	Open University
Substance	M	T	W	T	F	-	-	Cambridge University
Sell	M	T	W	T	F	-	-	APU
Workshop	M	T	W	T	F	-	-	APU
Workshop	M	T	W	T	F	-	-	Cambridge Computer Lab
Workshop	M	T	W	T	F	-	-	APU
Workshop	M	T	W	T	F	-	-	APU

HERE AND AWAY, UPCOMING

1-11 Feb	Bob Anderson	Sick	Princeton
2-21 Feb	Osamu Tsutsui	Sea Point, New Dorps	
18-19 Feb	Philipp Thorne	Head of Research, C.C.	London
11-14 Feb	Charles Heath	Paris	
12-13 Feb	Bob Anderson	Paris	
14 Feb	Bob Anderson	Baltimore, USA	
19 May	Bob Anderson	University of Canterbury	London
4 Jun	Bob Anderson	London	
17 Jul	Mick Lanning	Victoria	

THIS WEEK'S EVENTS

10-11 Feb	Work - PARC Bob Anderson attending	
10-11 Feb	9:00	Work - Philip Thorne (Sussex) Host: Charles King
MOF 10 Feb	14:00	Work - Susan Clome (Sussex) Host: Charles King
	14:00	AK: Brown Meeting Host: Maggie Bridge, Mike Norman, Mick Lanning, Alan Sales, Alan Zeyher attending. Conference Room
	15:30	TBA
	18:30	TBA - Paul Downhill 'Title' Hospitality and Hospitality at CRLW Design of Site Fall Three Services Home, London
11-13 Feb		Event: Basic Research HCl (open project meeting, Paris - Alan McLennan and William Jackson, Alan Zeyher attending)
11-14 Feb		Event: W. Jackson, Paris - Charles King participating
TUE 11 Feb		Work - Rob Onishi (Open University) Host: James MacKerel, P.S.
11-13 Feb		Work - Susan Clome (Sussex) Host: Charles King
	13:00	Work - Jim Lewis (OIC) Host: Mick Lanning
	15:30	TBA
	18:30	Serice Meeting on Meeting, Paris, France
WED 12 Feb	13:00	Work - Marc de Dalmé (OIC Holland) Host: William Norman
	13:00	Work - Owen Kaban (Dutch) Host: William Norman
	14:00	Work - Ian Lally, Mark Jarrett (Sussex) Host: Jim Downhill, P.S.
	14:30	Work - Tony Oller (Sussex) Host: Charles King
	15:00	Open House Meeting Host: Mike Miller, Sam Widdow, Kathleen Longhouse, Charles King attending
	15:30	TBA
13-15 Feb		Serice Meeting on Meeting, Paris, France - Bob Anderson attending
THU 13 Feb	9:00	Work - Philipp Thorne (Sussex) Host: Kathleen Longhouse
	11:00	Meeting - Paul Jackson, re: Image Research at Marine Centre, Cambridge University Architecture Department Organized. Kathy Carter attending
	12:00	Meeting - re: WAYS, Ian Mac, Albert Egan, Alan Brown (at IT Maths) with 3 guests (Lali) Richard Harper, Woody Marley, Richard Harper, Mike Miller attending
	15:30	TBA
FRI 14 Feb		Meeting - Herbert Dreyfus (University of California, Berkeley) Bob Anderson, Charles Heath attending
	15:30	Cambridge University Computer Lab meeting Open on Robotics and from various, including Leitch Thorne, New Materials Site. Kathy Carter attending
	15:30	CREAM TBA
SAT 5 Feb		
SUN 16 Feb		

UPCOMING EVENTS

18-19 Feb	Work - David Overstreich (University of Nottingham) Host: Paul Leff	
18 Feb	10:00	Work - Tony (Sussex) Host: Kathleen Longhouse
	15:00	Tue. Pub Meeting at Penetration
19 Feb	10:00	Work - Paul Bence (Computer International) Host: Alan Zeyher, P.S.
	15:00	Work - Paul Bence (Computer International) Host: Alan Zeyher, P.S.
20 Feb		Work - Paul Bence (Computer International) Host: Alan Zeyher, P.S.
24 Feb		Work - Paul Bence (Computer International) Host: Alan Zeyher, P.S.
25 Feb		Work - Paul Bence (Computer International) Host: Alan Zeyher, P.S.
26 Feb	10:00	Management Meeting Host: Bob Anderson, Alan McLennan, William Norman, Mike Miller, Charles Heath, Sam Widdow, Charles King attending
	15:00	Lab Meeting - Computer Lab - Bob Anderson attending
27 Feb		Work - Tony (Sussex) Host: Charles King
	15:30	Work - Tony (Sussex) Host: Charles King
	15:30	Work - Tony (Sussex) Host: Charles King
	15:30	Work - Tony (Sussex) Host: Charles King
28 Feb	TBA	Work - Tony (Sussex) Host: Charles King
4 Mar	TBA	Work - Tony (Sussex) Host: Charles King
11 Mar	11:00	Work - Tony (Sussex) Host: Charles King
	12:30	Work - Tony (Sussex) Host: Charles King
15 Mar	15:00	Work - Tony (Sussex) Host: Charles King
	15:30	TBA
16 Mar	10:00	Work - Tony (Sussex) Host: Charles King
17 Mar	TBA	Work - Tony (Sussex) Host: Charles King
18-20 Mar		Work - Tony (Sussex) Host: Charles King
22 Mar		Work - Tony (Sussex) Host: Charles King
	12:30	Work - Tony (Sussex) Host: Charles King
23 Mar	TBA	Work - Tony (Sussex) Host: Charles King
26 Mar		Work - Tony (Sussex) Host: Charles King
29 Mar	11:00	Work - Tony (Sussex) Host: Charles King
	15:30	TBA

Figure 1: The Calendar has three sections. The first shows a summary of planned lab attendance for that week; the second lists the week's events; and the third lists upcoming events over a period of months.

THE CALENDAR IN USE

A number of related issues characterise the way in which the Calendar is used. Overarching all of these is the notion of the Calendar's "formality". The Calendar is a well-established "official" document and is seen as the formal repository of a number of items of critical information.

The Calendar Administrator

A single individual is responsible for coordinating the Calendar and the information it contains. She collects information from lab members, collates it in a single document, and distributes it once a week both by electronic mail and on paper. The paper distribution is to lab members working in our building; the electronic distribution list includes colleagues and affiliates at other sites.

For most lab members, the Calendar is a source of information about upcoming talks, meetings and other events, and also includes summaries of upcoming visitors or periods of absence. The administrator's relationship to the Calendar and the information it contains, however, is very different. Many of her duties within the laboratory depend critically on the information which the Calendar holds. So, while most people see the Calendar as a repository of information about ongoing and upcoming events, the administrator views it as a central organising focus for her day-to-day activities. For instance, a Calendar entry indicating a visitor also implies responsibility for organising parking spaces and perhaps accommodation. Based on data

that has passed through the Calendar, the administrator also must also produce statistics on visitors, presentations, etc. for the lab's bi-annual reports.

The Calendar administrator has organised the structure of the document to support her own practices and the various responsibilities which Calendar entries represent. Indeed, the distributed Calendar frequently includes annotations she makes indicating the state of these other activities associated with particular items of information. These annotations are effectively private; although they are distributed with the Calendar document, few readers are aware of them, and fewer still understand their meaning.

Inputting and Collecting Calendar Data

To perform her duties, then, it is critical that the administrator can collect the information she requires. While individuals sometimes spontaneously send her notifications of upcoming events via a special "Calendar" electronic mailbox, the questionnaire responses indicate that most messages are sent in response to an automatic weekly electronic mail reminder to all staff. Information requests are divided into three principal categories—next week's attendance plans, upcoming events, and upcoming visitors or absences. Should an individual fail to respond to this message, the administrator will often follow-up with a personal reminder to submit a Calendar entry. Finally, if no information is forthcoming, the Calendar will show a "default" entry for that individual and mark it as such¹.

Distributing the Calendar

The Calendar is typically 3–5 pages long; representative pages are shown in figure 1. Most questionnaire respondents refer to it at the time it appears in their physical or electronic mailbox. After this, many do not consult the Calendar again, and either dispose of it or file it somewhere. However, others keep it around and refer to it over the course of the week. Of these, few report consulting it more than once a day, and the rate is often less than this.

Calendar readers report that the long-term information it provides is particularly useful. Most describe scanning the record of long-term events such as the list of upcoming absences and visitors to the lab over a period of two or three months. The format of the Calendar makes this easy; the information is often summarised on a single page. The brief summary of next week's attendance is also frequently cited as a useful source of information.

Having gained some picture of the role of the Calendar in the day-to-day activities of the lab members, we will now go on to consider the ways in which Khronika performs similar functions.

KHRONIKA IN USE

The use of Khronika within the laboratory is much harder to characterise than the use of the Calendar. First, unlike the Calendar, it is not a single, identifiable and physical artifact. Second, it provides a range of interfaces for inputting information, as well as a number of ways of scanning through it. Third, browsing may be replaced altogether through information retrieval by active daemons, replacing the scanning activity with automatic reminders. The presence of daemons, in particular, makes usage characterisations much harder; some users, for instance, claimed never to use Khronika, but yet referred to the (daemon-generated) audio reminders which they received. Since Khronika's interfaces are *environmental* as well as screen-based, Khronika reminders become less associated with a particular interactive system.

The most commonly reported means of accessing events in Khronika, however, is the workstation interface shown in figure 2. This provides a "week-at-a-glance" view; spatial layout is used to show the time and duration of events, and colour-coding provides cues as to event types. Selecting an individual entry causes the interface to present more detailed information on that event.

Although Khronika allows events to be entered which are accessible only to specific individuals, the primary use of

1. Like the administrator's private annotations, the marking of default entries is relatively subtle and often goes unnoticed. This information, then, is primarily used by the administrator herself

Khronika, and the use for which it was designed, is as a repository of public information.

Entering Khronika Data

The entry of information into Khronika is not triggered by a specific event like the Calendar's weekly reminders. Instead, individuals contribute information as and when they feel it is appropriate. The system provides a number of different interfaces for data entry; indeed, one of these allows electronic mail submission similar to that used for the Calendar, so that information can be sent to both with a single message. Unlike the Calendar, which is distributed once a week, information entered in Khronika becomes available immediately.

As with the Calendar, most users indicate that information is entered by the individual most concerned with the event it describes, *e.g.* a visitor's "host", or the speaker at an internal talk. However, a few individuals who rely heavily on Khronika for their information needs, and have a higher investment in it, will frequently enter information which is actually "owned" by someone else. Principally, the motivation is that this improves the "completeness" of the database, and justifies continued reliance on Khronika as a primary source of information. The information added in this way may come from a variety of external sources, including the Calendar, electronic mail, official announcements or informal communication with colleagues.

These individuals, then, take on an *informal responsibility* for the information in Khronika, partly because their working practices demand it. In some ways, this parallels the way in which the administrator's working practices require a certain correctness and completeness of information in the paper Calendar, and her formal responsibility for that information. Although the reasons for it are different, and in the case of Khronika this responsibility is much less formal and more distributed, some of the effects are similar. Both serve to maintain database correctness and increase the reliability of (and hence the validity of reliance upon) the system.

Retrieving Khronika Data

As with input, output from Khronika is not based on one specific system or medium, but rather can arise from any number of interfaces and subsystems. There are three principal ways in which information is extracted from Khronika—through *browsing*, *searching* and *automatic reminders*.

Most Khronika users browse the database frequently, sometimes several times a day. This is generally through the weekly browser shown in figure 2. A version of this interface automatically updates the display as new events are entered, making quick scans particularly easy. Many users keep a browser icon on their screen, which can be accessed quickly and simply. So, for users whose activities are primarily

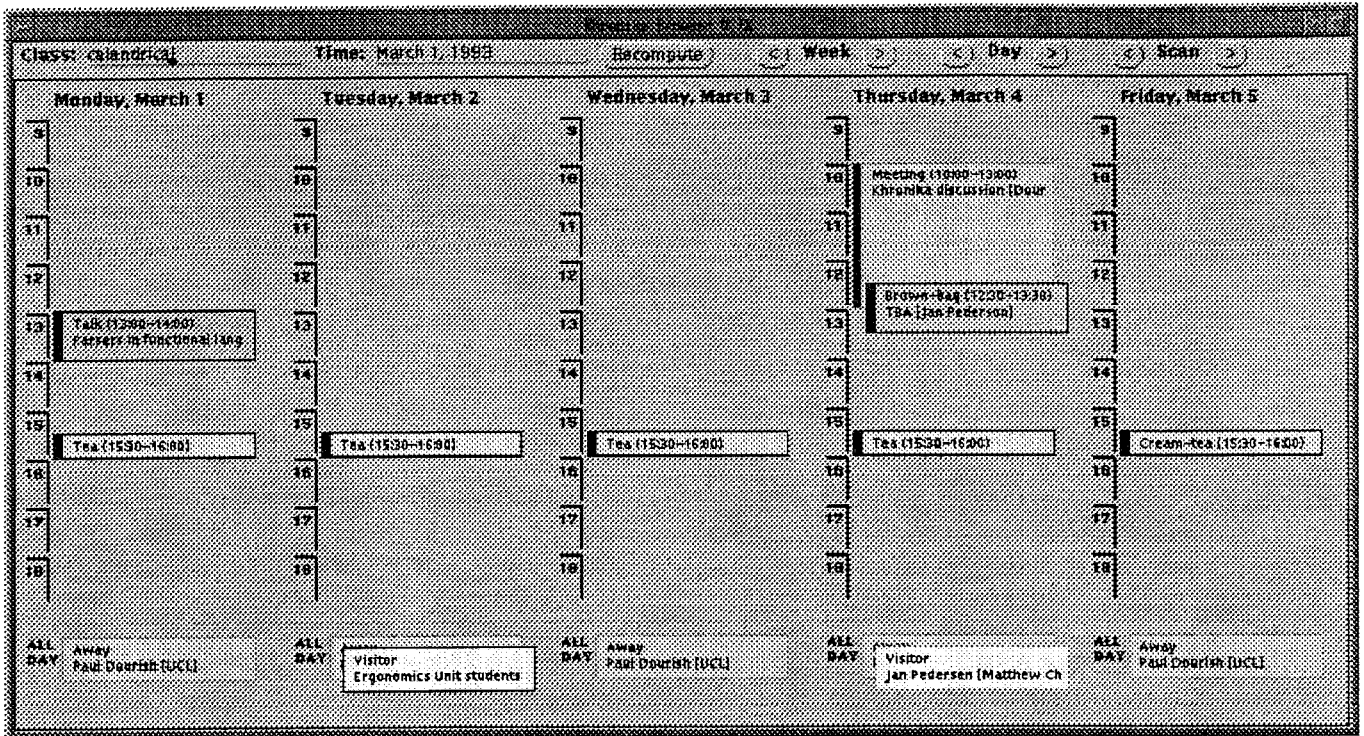


Figure 2: The most common interface to Khronika is the week-at-a-glance browser. Mousing on events will raise a pop-up window showing more details.

computer-based, it can be closer to hand than a printed Calendar.

Interface issues also affect patterns of searching in Khronika. (Other interface issues, concerned with information presentation, will be discussed later.) Searching Khronika for specific information is much less common than searching the Calendar. We have identified three principal reasons for this. First, the browser interface does not offer facilities for automatically finding specific pieces of information, resulting in much less searching activity. Second, it presents a short-term (weekly) view of the database, rather than the long-term overview given by the Calendar. The questionnaire indicates that Khronika is generally regarded as being more useful for short-term information. Third, questionnaire responses indicate that people browse Khronika much more frequently than the Calendar. This reduces the need to scan for specific details; this is essentially the *opportunistic rehearsal* noted by Payne [11].

Once information has been located, though, accessing the details is easy; selecting an event in the week browser will bring up a more detailed description of the event, as well as showing who entered the event into the database, and when. The Calendar does not offer this two-level browsing interface.

The majority of information output from Khronika, however, is in the form of automatic reminders generated by daemons. Khronika generates reminder events from

daemons as soon as the trigger event is entered. This means that, at any given time, there are a large number of events, spawned by daemons, waiting to be delivered at the time specified by the user in the daemon; and at the time of writing, the system currently contains almost 200 daemons, which have generated over 3500 pending reminders. The vast majority of reminders take the form of non-speech audio signals which are delivered to an individual's office via the media space. These signals are of particular use for short-term reminders (*e.g.* a meeting about to start; a seminar in ten minutes' time). Again, the largely short-term nature of Khronika use is apparent; this is a significant feature of people's use of Khronika.

COMPARING USAGE

Before going on to consider the components of contextual information which are being employed by the users of these two systems and the design implications which we can derive from these, it's worthwhile comparing the usage patterns directly.

Although Khronika and the Calendar manage similar information, their use is very different. The character of the information itself, and the processes of contributing, retrieving and interpreting it, vary quite widely between them. This raises a number of important questions. What factors do individuals use when deciding to use one system or the other? When the basic information is the same, what

factors cause people to interpret it differently in the different systems?

Certainly, some of these factors are traditional concerns of ease-of-use and appropriateness of each system for the working practices of the groups using them. However, we believe that there are other factors at work. In particular, we wish to deal with the way in which the systems present their information within a context which guides interpretation of the information and of the systems themselves.

In this section, we will look more closely at some of the differences between Khronika and the Calendar. It is convenient to be able to break these down into manageable units, so, although it is clear that they are strongly interrelated, we will deal with the *entry* and *retrieval* of information separately.

Entering Information

The value of each system is critically dependent on the completeness and accuracy of the information it contains. As information is received and generated within the laboratory, each individual has choices on how to distribute it.

A number of contextual factors play a part in this decision process. The organisational status of the Calendar, for instance, demands that certain information be put there; users comment that they enter information because they “feel obliged”. This status derives not only from the Calendar’s position as the official record of lab activities (*e.g.* vacations, visitors and talks), but also from understanding on the part of each lab member has of the Calendar’s audience and their view of the information it contains. This understanding includes an appreciation of both the nature of the audience, and of the role of the Calendar in their activities. The official status of the Calendar is reinforced by the weekly “calendar submission” reminders sent out by electronic mail, and is also personified in the administrator, so that an individual’s thought is often not “I’d better put this in the Calendar”, but rather, “I’d better tell Colleen”.

Khronika also has a notion of audience and role, albeit less formal. First, the audience is smaller and more local; unlike the Calendar, Khronika information is generally not accessible outside the laboratory. Second, Khronika information is not used to generate reports and statistics—indeed, it is not archived at all. Finally, information in Khronika can be amended or deleted much more easily than information in the Calendar (and Khronika will ensure that reminders are updated appropriately). All of these factors mean that less formal information can be entered in Khronika than in the Calendar; inappropriate dissemination of information is much less likely.

In selecting the appropriate channel for information delivery, then, the status, relevance and importance of the information has to be evaluated in the context of each system. Contextual factors such as the medium and timeliness of distribution, the organisational status of the system, and the perceived audiences of the systems, all play critical roles in this evaluative process.

Retrieving Information

The wider context of use is also critical when people retrieve information from either of these systems. This manifests itself primarily in two ways: the choice of which system to consult for some piece of information, and the process of interpreting information found there.

For those users who report using both systems, there is generally a distinction between systems based on the kind of information being retrieved. The choice of which system to refer to for some particular piece of information depends on a number of factors. Again, organisational status, timeliness and the nature of the medium are clearly important.

A major factor in interpreting information from the Calendar is the user’s understanding of the processes by which information is entered and coordinated in the Calendar. The presence of an item in the Calendar implies that the administrator has seen and processed it. Further, because the Calendar serves as an archival record, a certain status is accorded to aspects of the information it records. However, the Calendar is only distributed once a week, so information can only be seen to be “correct at the time of printing”.

Information in Khronika can be much more up-to-date, and there is no intermediate “administrator” separating information entry from database storage. However, because Khronika lacks the official status of the Calendar, it is often incomplete, and this is one of the most common problems quoted by questionnaire respondents.

Khronika annotates information with records of when the information was entered, by whom, and when it was changed. When combined with knowledge about the organisational responsibilities and work practices of colleagues, these facts are critical in the interpretation of the information which is retrieved. For example, the knowledge of the Calendar administrator’s other responsibilities, such as booking accommodation, might lead us to decide that, in case of conflict, the Calendar is more likely to contain the correct information about the arrival date of a foreign visitor who is to give a talk. However, the administrator does not typically attend internal seminars, so Khronika might well be regarded as more likely to contain the correct title for that visitor’s seminar during her visit. Even then, knowing who entered the seminar announcement into Khronika might cause a reassessment; different trust levels would be assigned to the seminar coordinator, the visitor’s host, or

some other unconnected lab member. These might be further resolved by looking at the time of entry, and so forth.

Clearly, input and retrieval of information are not independent. An understanding of the process of information entry can also affect interpretation. An event in the Calendar carries with it an implicit statement that it is “definite” and “settled” (unless otherwise marked). Partly, this is due to the Calendar’s official status; but it also relates to the interaction with an individual who mediates between information sources and the Calendar document. This means that repeated changes to Calendar information are less likely than those to Khronika information. Events in Khronika may be much more tentative; partly because it’s easy to change information later, and partly because Khronika arranges that information will primarily be extracted (by daemons) only in response to specific registered interests of individuals. There is less worry that entering and later changing some tentative information in Khronika is going to “bother someone”.

Presentational Issues

It is clear from reports of usage that one important factor in the differences between Khronika and the Calendar is the nature of the presentation of essentially similar information. While the basic data which the systems present are the same, the nature of the presentation makes different information visible. This is a generic issue in information presentation [5].

Consider the information as it is presented in the Calendar and Khronika (figures 1 and 2). The Calendar layout separates information according to particular events, as opposed to particular times, with the result that conflicts between events (*e.g.* a meeting scheduled at the same time as a seminar) are hidden in the details of the information. The standard Khronika browser, on the other hand, uses a spatial layout based on event times and durations, and hence highlights this information. In looking at the Khronika browser in figure 2, it is immediately obvious that there is a conflict between the meeting and the “Brown Bag” seminar on Thursday.

The Calendar, on the other hand, presents in its third section a compact summary of upcoming events and visits over the next few months. This information is much less accessible within Khronika; in the standard browser, it requires paging forward over entries for weeks at a time, and in other browsers it requires the construction of an appropriate query.

It seems clear that certain aspects of comparative use of the two systems, such as a focus on immediate or short-term events in Khronika and more longer-term information-gathering are to do with the presentational issues which are an implicit part of their interfaces; indeed, the use of non-speech audio reminders triggered by daemons in Khronika is simply a presentation mechanism which is impossible with

the Calendar. However, these factors are primarily concerned with *retrieval*; they do not deny the importance of contextual cues which cause the same information to be *interpreted* in different ways once retrieved.

COMPONENTS OF CONTEXT

We have reviewed the way in which Khronika and the Calendar are used by a specific work group. One of the most striking observations is that the use of these two shared information systems goes beyond the naive view of information retrieval. Instead, around the systems, users have built up a more elaborate set of practices which concern not simply retrieval but also evaluation of information, and so on. These practices are enabled by a variety of resources, which together we have classed as *context*.

Both Khronika and the Calendar provide context for any information which they hold. They do this in very different ways. In the Calendar, this context comes from its processing by the administrator. Khronika, on the other hand, tags information with parts of the context of entry *e.g.* the name of the person who entered it, and the date of entry or modification. This can be used in judging the reliability of the entry.

Many of the resources which make up this context are not derived from the information itself, but are based on a body of largely tacit organisational knowledge. So we cannot *enumerate* the resources and fully describe the information context as it is used. However, in looking at the usage of these two systems we can begin to identify some specific resources which are used to manage the information which the systems contain. These include:

1. *Ownership and responsibility.* We have seen how both Khronika and the Calendar present information with contextual cues which indicate ownership of information, and responsibility over its correctness. These are used by individuals in determining correctness and appropriateness of information.
2. *Medium and mutability.* The medium in which information is stored and distributed plays a major contextual role in information entry, retrieval and interpretation. While convenience and ease-of-use are important factors, this relationship to medium goes deeper. A more mutable medium, for instance, is more likely to carry tentative information, since there is less commitment being made in submitting it; similarly, this knowledge will affect the interpretation of information retrieved. Many of these properties can be analysed from the perspective of ecological affordance theory [3].

3. *Timeliness*. Partly related to the issue of the medium is that of timeliness. Our two systems have been associated with particular mechanisms of information collection—either individual entry from a workstation, or weekly collation by an individual. These processes have different timescales, and this forms part of the context in which individuals make decisions. It affects information contribution (since an individual may decide that information would be delivered too late), and information interpretation (e.g. when information is clearly out-of-date).
4. *Organisational status and relevance*. Different kinds of information have differing degrees of relevance to individuals and to the organisation in which they operate. Similarly, the various information systems which might be available also carry with them some notion of organisational status and relevance for particular functions. This again is an important consideration for individuals both in contributing and interpreting data.

These are *components* of the information context, not orthogonal dimensions. Their interactions are complex. For instance, considering the example above concerning the decision against submitting information to a less timely medium, an understanding of organisational relevance might mean that the information will still be submitted. Indeed, this happens in our systems; because the Calendar is used to generate six-monthly statistics, it is important that information be submitted even when it is no longer worth distributing.

DESIGNING FOR INTERPRETATION

Our studies, then, have shown the way in which the uses of Khronika and the Calendar differ, despite the similar information they hold. They have pointed to the fact the information management in such systems is subject to a set of complex and subtle practices which evolve over time. These practices are enabled by a variety of resources available within the information management systems which together provide a context for the information which can be used to assist the management process and to interpret the information.

The context is an implicit component of all information which we receive, and we are all quite used to managing it in everyday life. This is illustrated by the example given in the introduction of the way in which our judgement of the reliability of a telephone number will differ depending on the source and form of the information as it appears to us. Since actual, day-to-day information management is enabled by, and conducted through, the use of these resources, then it becomes clear that, if electronic information systems are to be as flexible and valuable as non-electronic equivalents,

provision must be made for such resources to be provided, and for the retrieval process to involve them. In shared information systems, where many different people contribute to and use the information base, provision for such context is even more critical. Failure to adequately support such cues, which underpin normal, everyday information processing, is a known source of problems in collaborative systems [7].

In trying to create systems which acknowledge the role of contextual information, some serious associated concerns arise. The first is that the use of contextual information emerges from the use of the system itself, and from the social setting in which it is based. Aspects of this use of context, then, are *coadaptive* [9], as the use of the system and the interpretation of the information grow up around each other, and affect each other bidirectionally. Further, the context extends outside the computer system itself, and can also include other systems, practices and organisational factors which exist in the environment or the workplace. As we've seen in the case of Khronika and the Calendar, these interactions may be subtle and indirect, but yet, they remain importantly interdependent. These issues are explored in greater detail by Markus and Connolly [10].

So, the issues which surround the design of information systems which can provide the resources needed for flexible use are complex. We can summarise some of the main design implications from our studies:

- Contextual cues are important resources for information management. The failure to exploit such cues, then, is likely to lead to less usable and less flexible systems. It is important that the significance of these cues is recognised, and that electronic information systems attempt to provide them. The design process for information systems should be sensitive to the nature of the decisions made based on the information they supply to users, and to the factors which influence this decision-making process.
- The practices which evolve from this use of contextual cues are informal, implicit and evolving. Hence, it is generally not practical to provide, within the system, actions which are based on contextual information. The contextual information is of value to the end-user in interpreting the information, rather than to the system in making inferences, even on the user's behalf. This leads to a model of shared *information*, with the system serving as a focus for information and browsing, rather than shared *action*, where the system performs tasks which are traditionally performed by individuals. We believe this to be one of the major distinctions between Khronika and that class of shared calendars which attempt, for instance, automatic meeting scheduling, and it is to this that we attribute much of Khronika's success.

We would argue that this distinction is an important one to make in the design process; and it is one which follows quite directly from an appreciation of the value of informal contextual information in decision-making.

- Contextual information is frequently low-level information which is it easy for a system to collect (such as names, entry times and so forth). This information can then be recorded opportunistically along with the system data, imposing little additional overhead. Information systems, then, should collect this information when possible and record it along with the data for later presentation.
- Contextual information is most useful in *browsing*—indeed, the availability of such information makes browsing much more productive. Automatically-generated reminders, for instance, are focussed much more on the actual data rather than the context which surrounds it. The value of the context is earlier in the retrieval process. It is important, then, that browsing activities be supported and that contextual information be presented at this stage.

So, in designing information systems, we cannot consider the information which will be managed by a given system without considering questions such as:

1. What contextual cues should be provided in order to assist users in interpreting the information correctly? The information, in and of itself, is of little use if the interpretations cannot be made.
2. How will the appropriate set of interpretive practices develop, and how can this process be seeded and guided?

We have given some guidelines for the ways in which contextual information can be collected and used in shared information systems. More investigation is necessary in order to reveal more intricate relationships between the information which these systems carry and the context in which it is embedded. These are complex issues, but they are at the heart of the design of large shared systems. In looking at the use of context in current systems, and some of the components which are employed in the interpretation of information, we can begin to address them by making explicit context a part of our development model.

SUMMARY

We have presented some early findings from our studies of two shared information systems. This work was an excellent opportunity to study two systems dealing with similar information in the same organisation, and hence to look at the differences in their use.

We have shown how the differences in use appear to be based primarily not on the information itself, but on a variety of contextual factors which act as resources in the process of *interpreting* the information. This interpretation is critical—if information is to be of use to an individual, it must be perceived to be correct and pertinent. The resources take a variety of forms, sometimes technical, sometimes social and sometimes organisational, and they all contribute to the interpretation of the information.

Clearly, as we design electronic information systems for larger groups and more dispersed organisations, we must be sensitive to these factors. Information *retrieval*, in its purest sense, is simply part of the challenge—once the information has been retrieved, it must be suitably contextualised in order to be used. As designers of these systems, then, it is important that we have an understanding of the contextual factors involved, which may be unique to particular organisations or environments, and that we investigate the ways by which our systems can support contextual interpretation of information. Such support can make the difference between a system which merely provides multiple access to information stores, and a system which is an effective support to organisational work.

Acknowledgments

Tom Moran instituted the Calendar and first proposed an electronic equivalent. Lennart Löfstrand designed and implemented Khronika. We thank our questionnaire respondents, interviewees and particularly Kathleen Loughnane, for helping us understand the relationships between these two system, and Peter Brown, Bill Gaver and Allan MacLean for helpful comments on an earlier draft of this paper.

References

- [1] Beard, D., Palaniappan, M., Humm, A., Banks, D., Nair, A. and Shan, Y., "A Visual Calendar for Scheduling Group Meetings", Proc. CSCW'90 Computer-Supported Cooperative Work (Los Angeles, Ca., October 1990).
- [2] Gaver, W., "Sound Support for Collaboration", Proc. ECSCW '91 Computer-Supported Cooperative Work (Amsterdam, Netherlands, September 1991).
- [3] Gaver, W., "Technology Affordances", Proc. CHI '91 Human Factors in Computing Systems (New Orleans, Lo., April 1991).
- [4] Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. and Buxton, W., "Realising a Video Environment: EuroPARC's RAVE System", Proc. CHI '92 Human Factors in Computing Systems (Monterey, Ca., May 1992).

- [5] Gilmore, D., "Visibility: A Dimensional Analysis", in Daiper and Hammond (Eds.), *"People and Computers VI: Usability Now"*, Cambridge University Press, Cambridge, 1991.
- [6] Greif, I. and Sarin, S., "Data Sharing in Group Work", Proc. CSCW '86 Computer-Supported Cooperative Work (Austin, Texas, December, 1986).
- [7] Grudin, J., "Why CSCW Applications Fail: Problems in the Design and Evaluation of Organisational Interfaces", Proc. CSCW '88 Computer-Supported Cooperative Work (Portland, Or., September 1988).
- [8] Lövstrand, L., "Being Selectively Aware with the Khronika System", Proc. ECSCW '91 Computer-Supported Cooperative Work, (Amsterdam, Netherlands, September 1991).
- [9] Mackay, W., "Users and Customisable Software: A Co-Adaptive Phenomenon", PhD Thesis, Sloan School of Management, MIT (Cambridge, Mass., 1990).
- [10] Markus, L. and Connolly, T., "Why CSCW Applications Fail: Problems in the Adoption of Interdependent Work Tools", Proc. CSCW '90 Computer-Supported Cooperative Work, (Los Angeles, Ca., October 1990).
- [11] Payne, S., "Understanding Calendar Use", *Human-Computer Interaction* 8(2), 1993 (in press).