Making Contact: Getting the Group Communicating with Groupware

Andrew Cockburn

Department of Computer Science
University of Canterbury
Christchurch, New Zealand
agc@cosc.canterbury.ac.nz
+64-3-364-2774

Saul Greenberg

Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
saul@cpsc.ucalgary.ca
+1-403-220-6087

ABSTRACT

While groupware is readily available, people on wide area networks—such as the Internet—have considerable trouble contacting each other and setting up groupware connections. To pinpoint why this occurs, this paper identifies human factors critical to getting a group communicating through groupware. It addresses how people find suitable partners, and how people choose appropriate communication mediums. These factors are discussed in detail, and form a design foundation for systems that promote social presence and that integrate communication.

Existing systems are critically reviewed and shown to be inadequate for general use over a wide area net, for they either do not meet some basic design criteria, or they require a very high technological entry level that is beyond the reach of most computer users. As an alternative, the paper presents the design considerations behind TELEFREEK, a flexible, extensible, and customizable platform for collaboration. Drawing on resources freely available to the Internet community, TELEFREEK assists people making contact with others, and integrates access to common communication facilities.

KEYWORDS

contact facilitation, casual interaction, coordination, computer supported cooperative work, groupware.

INTRODUCTION

Groupware developers are creating special purpose systems supporting communication, coordination, and collaborative work. We now have sophisticated mailing systems, coordination-based tools, video and media spaces, brainstorming tools, collaborative drawing and authoring systems, and so on. Enthusiasts notice and applaud the many small-scale groupware successes, even though most are restricted to either local installations or tightly-connected distant sites. The truth is that we still have a long way to go before a large community connected by a wide area net can be served effectively by groupware. The two issues this paper concentrates on are that people have trouble getting in touch with one another, and that it is hard for them to choose and establish an appropriate communication or groupware channel.

As with physical interactions, an essential prerequisite for geographically dispersed groups is that people must be in contact with one another. Relationships must be established, and many interactions—including informal ones—may be required for people to find partners [22,31]. Yet the bottleneck to rich spontaneous interactions is distance [22], and users of wide area networks will be at a disadvantage unless a prosthesis that overcomes distance barriers is available. Many mundane factors interfere with making contact. People must know electronic addresses and even machine names. People must ready software, equipment, and each other well in advance for real-time remote conferencing. With video conferencing and media spaces, people must be in the (usually few) rooms that have the media equipment available. For informal interaction, people must find each other with minimal effort.

People must also select one or more of the many communication channels and applications that may be available to the group. This can be a difficult task. From a technical perspective, sites may not have the same software; workstations may not support the necessary media (such as digital audio); specialized equipment may not be available (such as video cameras and microphones); low bandwidth and high latency may limit interactions, and so on. From a human perspective, the communication channel or groupware must match a group's task (e.g. real time text editing), and accommodate how people are available (e.g. asynchronous vs. real-time). If people cannot make contact through an appropriate medium, then groupware systems—no matter how many are available or how eloquent they are—cannot be used.

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

Petitioners
Exh-1015, p. 001

While few groupware applications facilitate making contact, technology does have potential to bring distance-separated people together. This paper examines three general issues in communication:

- how people find suitable partners
- how people establish contact
- how people select their communication channels.

In particular, we focus on relatively low bandwidth networks that are heavily populated, such as the Internet. Because people on these networks have fairly limited access to resources, we want to build on facilities that already exist on the network. We also want to show that a fairly simple system can meet many of the requirements now handled only by high bandwidth and technology-intensive systems.

The paper begins by listing critical human factors for systems that help people contact one another. We believe that the factors raised form the minimum design foundation for any infrastructure supporting groupware. The next section reviews current systems that facilitate social browsing and directed encounters, and discusses why they are not yet suitable for the Internet and its community. Finally, we will describe the design premises and current implementation status of our "work in progress" system. TELEFREEK is an integrated communication and social presence assistant that draws on resources freely available to the Internet community.

DETERMINING THE WHO AND HOW IN COMMUNICATION

Several factors commonly affect the efficiency of cooperative work. The first problem is getting in touch with the right people. Specific needs of communication initiators are thwarted by inadequate knowledge of suitable recipients. They often need to know who is around, how they can be contacted (synchronously and asynchronously), who is pertinent, whether they are available for interruption, and whether the social status of the recipient is appropriate for communication to occur.

This section describes the major human factors affecting the user's decision in selecting particular people and mechanisms for communication, as summarized in Table 1.

Getting in touch with the right people

Social presence plays a dominant role in helping people form and maintains working relationships. Social presence is concerned with the whereabouts and availability of potential communicants [22]. We keep in touch with the people around us to nurture relationships. Accidental and intended contact through real-time communication is what we regularly do; this is why we wander hallways, go to public places, and phone people. Yet intentional real-time contact is the hardest to achieve. Troublesome and annoying difficulties, such as telephone tag, can make users resort to less rapid but more reliable mechanisms—post-it notes stuck on a desk, email, and even fax. While responses are not immediate, receipt of the message is virtually guaranteed (although now the recipient has to find the originator).

1. Getting in touch with the right people

a) Social presence plays a dominant role in helping people form and maintain working relationships.
⇒Computers could give its users a sense of who is reachable through the network and who is available at the moment.

b) People need to find pertinent members of the community.
⇒Computers can direct users with specific needs to the individuals most likely to be able to satisfy their needs.

c) People need to find particular people.
⇒Computers should make it simple to contact a specific person.

d) The social status of participants may affect who is allowed to know about others, and who can initiate communication.
⇒Computers should make people aware of the social status of potential contacts.

e) The recipient may not want to be disturbed by the caller.
⇒Computers could allow recipients the choice of accepting or rejecting a call. Alternately, they could track one's interruptability status and either notify others of it, or disallow access.

2. Choosing the right communication channel

a) The task strongly influences the minimally acceptable communication channel.
⇒Groupware should allow people to choose an appropriate communication medium.

b) The imbalances of available facilities restricts how people can communicate with each other.
⇒Computers can tell call initiators what equipment and communications channels are available for a particular meeting.

c) The intended period of interaction affects our choice of a communication channel.
⇒Computers should make users aware of the broad range of available communication mediums and their delay characteristics, and should encourage media switching when necessary.

d) People must overcome the inertia involved in switching to better communication channels.
⇒Computers should allow users to swap easily between communication channels, and should take an active role in easing the communicants' transitions between channels.

e) People expect conversations to continue across time and space boundaries
⇒Computers should allow communications to migrate naturally between boundaries. Distinctions made between asynchronous/synchronous and co-located/dispersed groupware should be minimized.

Table 1. A summary of the human factors affecting ways groups get together for cooperative work, and how computers can provide support.

Petitioners
Exh-1015, p. 002

Before the advent of networked computers, access to social presence information was limited to means such as walkabouts, clock-in cards (stating where people might be, not necessarily where they are), or by actually establishing contact through accidental and pre-planned encounters. Networked computers, however, can readily provide a variety of information about the community of users. We can easily imagine that our computers can pass on our activity level, our availability for interruption, how we can be reached, what communications media we have available, and so on. Other modern technologies, in conjunction with computerized techniques, allow closer and more continual updating on the whereabouts and activities of individuals. Such systems, which include media spaces and "active badge" technologies, will be discussed later.

Because social awareness and eased access to potential collaborators are important for developing working partnerships, groupware should give its users a sense of who is reachable through the network and who is available at the moment, and should allow easy access to particular people.

People need to find pertinent members of the community. When people need help solving particular problems, they have to find the right people to ask. There is usually a sense of urgency, so timeliness of response is important. While blanket addressing may resolve their problem quickly, there is substantial cost in the form of wasted time and interruption to the majority of those reviewing the request. Even when we do have a particular person in mind, they may not be available, and alternatives must be pursued.

While general mechanisms supporting social awareness will help, computers can assist more aggressively by directing users with specific needs to the individuals most likely to be able to satisfy those needs. Schemes for finding pertinent people include: databases of information about network users—what they do and what groups they belong to; filters providing awareness views of only relevant experts [23]; and expert systems that automatically answer queries previously encountered, and finds appropriate experts for new ones [1].

People need to find particular people. The selection of communicants will often be pre-determined by a user's particular tasks and needs. The more specific the need, the more likely that work can continue only with a specific partner. For example, writing a paper will involve tight collaboration between co-authors, and there may be no need to contact other people in the community at large. Because people may need to contact their partners, groupware should facilitate finding specific people.

The social status of participants may affect who is allowed to know about others, and who can initiate communication. The relative social status of communicants can raise complications stemming from protocols (formal or informal) for communication through organizational hierarchies. While it may be acceptable for managers to phone subordinates, the reverse need not be true. When it could be an issue, groupware should make people aware of the social status of potential contacts.

The recipient may not want to be disturbed by the caller. Attempts to establish communication commonly cause unwanted interruptions. This situation is typically unsatisfactory to all parties—the caller is sorry for the interruption, while the recipient must deal with the interruption and afterwards recover their previous train of thought or action. Personal assistants provide a manual solution to this problem for those who can afford to employ them, but this is rarely the norm for computer communication.

Computers can offer a more accessible alternative. First, the computer should allow the choice of accepting or rejecting a call, as telephones do. Of course, this could be a problem if the sender knows the receiver is there! Similarly, the computer can allow one to screen calls surreptitiously, as we now do when we hear someone leave a message on our answering machine, only picking up the phone if we really want to talk to them. Second, users can record their interrupt status. Senders can review this, and use their own judgment to decide if the call should be placed. Alternatively, the system can enforce an interruption structure that will control whether or not others are allowed to complete a call. While there could be difficulties with any of these methods, we would expect new social protocols to evolve that make these approaches acceptable to different groups.

Choosing the right communication channel

The task strongly influences the minimally acceptable communication channel. The collaborative task will, to a large extent, dictate appropriate communication mechanisms for information exchange. For example, if the task is transferring a textual document to another person, email would be preferred over telephone dictation. For real-time collaborative editing, a group editor with a voice link would be preferred over each partner working on their own document and sending each other periodic updates. While the other channels are usable, they will be a bottleneck to the work.

Of course, it would be simple to ask for the highest quality communication channels available on a network, for these would certainly be capable of satisfying the users' requirements in rate of information exchange [22]. Yet this can be tremendously wasteful of resources. What we really need to know is the *minimum* bandwidth necessary for user and task satisfaction. Increasing the bandwidth of channels, while appealing, does not necessarily assist collaboration and is thus not cost-effective [12,13,33].

Because tasks often determine what medium is appropriate, groupware should allow people to choose an appropriate communication medium.

Petitioners
Exh-1015, p. 003

The imbalances of available facilities restricts how people can communicate with each other. Obviously, communication facilities are restricted to those available at the initiator's site. More problematical for the initiator, however, is knowing what facilities are available to the receiver. It is no good trying to establish a video-phone link if the other person does not have a video camera wired into the network!

Groupware supporting social presence information can help ease these problems. It can note what equipment is available, and can tell the initiator what the possible communication channels are. It can also provide information about alternate ways of contacting someone. For example, if the system sees that a colleague is away from the office, the relevant home or mobile telephone numbers could be offered.

The intended period of interaction affects our choice of a communication channel. The period of interaction is the cycle time related to various aspects of interaction. Before computers, the three most common means for interaction (aside from face to face) were the telephone for synchronous and immediate contact, memos for local asynchronous contact with quick turnaround, and surface mail for communication tolerating a turnaround time of several days. Technologies, such as facsimile machines and networked computers, have made several new channels of communication available, increasing both the choice of media and the range of supportable interaction periods. We are now at liberty to select a medium that satisfies our needs for responsiveness.

The period of interaction suggests several crucial factors affecting our selection of media.

Task period is the time over which each unit of task activity must be completed. An example is cycling paper revisions between people.

Environment period is the delay imposed by the environment. These include time-zones, and recognizes that it is rarely acceptable to telephone someone at 3AM.

Propagation period is the delay imposed by the communication technology and how often people check that channel for activity. For example, this period is small for telephone interaction and high for surface mail.

Perceptive period is the maximum delay that users feel is tolerable for successful interaction. With real-time communication, a three second propagation delay on a satellite audio link is likely to disrupt conversation and therefore be equally unacceptable to all conversants. When using asynchronous technology (such as e-mail), there could be a conflict in the perceived period. While one correspondent may be content to take two weeks responding, the delay could frustrate others.

Because the duration of these periods is highly variable and subject to exceptions, it is perhaps unreasonable to expect any particular groupware to establish a balance between interaction delays *imposed* by the environment and support technology against those *desired* by correspondents. A reasonable alternative is for the computer to make users aware of the broad range of available communication mediums and their delay characteristics. Because the desired period of interaction can change over time, the computer can also encourage media switching when the conversation is not progressing as well as it could.

People must overcome the inertia involved in switching to better communication channels. The potential for fluid switching between communication media deserves noting as an important factor affecting the choice of mechanisms for interaction. In many circumstances, messages received on one device (such as a telephone) should stimulate further messages using a different and more appropriate mode of interaction (such as email). Yet there is a tendency for people to stay on one communication channel, which can result in inappropriate communication mechanisms being maintained for no express purpose. We call people's ability to ignore other communication opportunities "media blindness." This phenomenon is exemplified by several unsatisfactory email exchanges that finally results in a query "What's your telephone number?" The problem of media blindness is exacerbated by the lack of integration of most communication systems. As Bair [3] asks:

> "How can we extend electronic media to meet user's needs?... we know that the use of a variety of integrated media, selecting each for the appropriate purpose, is the ideal situation."

Implications for computer systems from this observation are twofold. First, systems should be flexible enough to allow users to swap easily between communication channels. Second, systems could take an active role in easing the communicants' transitions between communication channels. For example, a person receiving an email message may be offered a variety of appropriate ways to reply, which could include email, fax, phone, video links, specific groupware applications, and so on.

People expect conversations to continue across both time and space boundaries. A common classification of groupware uses two dichotomies: the first based on whether the group is co-located or dispersed; the second distinguishes between synchronous or asynchronous communication [20]. While these distinctions accurately categorize the explicit support provided by many groupware systems, they provide no insight to the actual communication requirements of system users [2]. For example, channels intended for synchronous use are frequently applied in an asynchronous manner. One case is how the telephone has been augmented by an answering machine; another is a real-time group sketchpad used as a post-it note [16]. Similarly, asynchronous mechanisms can be used synchronously, as we often see in the rapid exchange of email messages. The co-located/dispersed distinction is also a questionable one. Systems for remote conferencing (such as group editors) should—in principle—be quite usable for co-located people working together.

Petitioners Exh-1015, p. 004

The dangers of these distinctions lie in promoting the notion that systems necessarily support only one quadrant of these communication styles. Resulting systems will likely be difficult or impossible to use if people's needs shift to the other quadrants. Yet computers and modern telecommunication facilities can adjust to the pace and place of interaction: asynchronous media such as email can now be transferred locally in seconds and globally within minutes. With high-speed modems and wide-area networks, the distinctions between local and remote computing is becoming increasingly nebulous.

Designers must be aware that their groupware will be used in ways they had not foreseen, and supporting or easing transitions between communication methods will enhance a system's perceived value. Indeed, the better the system the more likely it is to be used outside its planned domain. As Ellis, Gibbs and Rein [10] state "a comprehensive groupware system might best serve the needs of all the quadrants."

RELATED WORK

While there are a variety of systems available for making contact, most fall under the categories of social browsing, and whereabouts and availability systems. This section critically examines a few of these systems. Where appropriate, we will relate their capabilities to the points raised in Table 1. We close by discussing why current approaches are poor fits to the Internet community.

Social browsing systems for casual interaction

Casual meetings and chance encounters are essential for spontaneous interaction. Because distance decreases the likelihood of these events, several systems have been developed to support "social browsing" over a network.

The simplest technical support for social browsing simply lists users available on some (or all) machines on the network. In UNIX, for example, the *rwho* command lists who is on the local network, what machines they are on, when they logged on, and their "idle" time i.e. how long it has been since they entered any input. Similarly, *rusers* allows one to poll remote machines. While knowing who is logged on does not guarantee that someone is actually at their machine, a short idle time gives a reasonable indication of presence (point 1a in Table 1), and where they are (point 1c).

Most other systems use video for social browsing [4]. Spontaneous "drop-in" interactions (point 1a) between people at distance sites can be encouraged by placing large video screens in common areas (such as coffee rooms). The first example was Xerox Parc's VIDEO WALL, where a slow-scan video connection was placed between two research laboratories located in California and Oregon [15,32]. In general, VIDEO WALL worked. Goodman and Abel [15] reported that VIDEO WALL was used for both spontaneous (70%) and planned (30%) interactions, and that people used it for both social and technical communications. Because people can see what others are doing, they also know if they can be interrupted (point 1e). Fish, Kraut and Chalfonte [11] discuss and analyze a similar system called VIDEOWINDOW.

While VIDEO WALL directly connects two physical locations, a "media space" attempts to create a virtual community where everyone has instance video access to everyone else from their offices [4]. Bellcore's CRUISER was designed on two premises: 1) users can browse a virtual world seeking social encounters (point 1a); and 2) users can construct, organize and populate the virtual world independent (within reason) of the physical world (point 1b,c) [12,31]. Users can also set privacy levels on how others can peek into their offices (point 1e). There are three methods for browsing.

- A *jump* supports a direct planned movement to a physical location. A user selects a remote location, and the image captured from the camera at that location appears on the screen.
- A *path* extends the jump idea by listing a sequence of locations and the order in which to visit them. This, in effect, becomes a "virtual hallway" through which the user can walk through.
- A *random walk* is similar to a path, except that CRUISER generates the locations either at random, or through a function satisfying some user desire.

Louie, Mantei and Sellen [25] present a taxonomy of six phases of human communication. One finds participants in the *pre-communication* phase, attracts their attention in the *attention* phase, and then uses the *greeting* phase as a ritual for negotiating how communication will proceed. The main communication occurs in the *maintenance* phase, followed by a highly structured agreement to end the session through a *closing* phase. The final *fade-out* ends the communication, and as the name indicates, it is rarely abrupt.

Within the above framework, the CAVECAT media space project [28] pays particular attention to the opening phases of making contact, and several metaphors have been designed to facilitate each phase [25]. In particular, they have developed person, time, and space centered views and metaphors for finding people in the pre-communication phase, and they list a variety of methods commonly employed in media spaces to attract attention of others.

Other experiments in media spaces are ongoing, e.g. RAVE [14], VOODOO [24], and TELEPRESENCE [5]. An excellent media space survey is found in Bly, Harrison and Irwin [4].

Whereabouts and availability systems for directed encounters

Natural compliments to social browsers are systems that enhance constant awareness of who is about, what they are doing, and whether they are available for conversation. The intent is to support one person trying to get in touch with another particular person.

One simple facility is the Unix *finger* command. When invoked with a person's login name, it returns: the person's

Petitioners
Exh-1015, p. 005

real name; if and where they are logged on; their idle time; and the last time they have read mail. It will also return the contents of a user-created file that typically contains one's title, contact address, a phone number, work plans, and so on (points 1b-d). More formally, the X.500 protocol provides a directory service that can be queried [30]. For example, people can ask the directory service to return the email address for "Saul Greenberg at Calgary".

MESSENGER extends a finger-like command with a means to establish contact with others [17]. It is a Macintosh desk accessory that shows a person a list of all others who have recently moved a mouse or touched a key. One or more active people may be selected, and a text message can then be sent to them as a post-it note (pop-up window). Message recipients can then view the message and reply.

The natural relationship between human requirements for social presence information and whereabouts information was observed during evaluation of CRUISER [12]. Users often established a video connection to an empty office, so that they could see when the missing occupant returned (point 1c). Surprisingly, such an "ambush" would usually result in one physically going to the other's office, rather than a videophone link.

Variations of this theme are systems that create an awareness of selected groups of people. PORTHOLES [9] uses low-bandwidth video technology to maintain a "sense of community" between both local and distributed sites. Cameras in each office periodically frame-grab a low-resolution image and transmit it to other sites; the "community" is presented on one's screen as an array of images (points 1a,c,e). While PORTHOLES can then be used to establish email and videophone contact with others (point 2d), the designers note that its major use was in locating colleagues.

Perhaps the most controversial system promoting awareness is the Active Badge LOCATOR, developed by Olivetti Research [34]. These are small badges worn by people that continually transmits one's location at the work site. A central database collects an identity code, the room location, and the current time. Applications can then make use of that information [19].

- Queries can be sent to the database asking about the current location of any individual (point 1c). Similarly, people can initiate point to point video and audio links to others by referring to them by name, rather than by address (point 2d).
- Location and movement data can be superimposed on a computerized floor plan.
- A personal diary can be built keeping information about daily activities (the PEPYS system [29]). Events such as meetings can be automatically noted along with their location.
- One can browse through events (past, present, and future), and "event daemons" can provide automatic notification of events within the database, such as meeting arrangements and other appointments (the KHRONICA system [26]).

While benign use of the LOCATOR appears promising, its possibility for abuse has raised concerns in the CSCW community [6]. As Harper, Lamming and Newman [18] note:

> "Locator information for one group of individuals in one particular context is a practical help, for another group that same information may open to question issues of power, control, entitlements and potential change in organizational circumstance."

Why current approaches are poor fits for the Internet community

Current approaches to making contact are not particularly appropriate for the general Internet community. First, the high technology adopted by most existing systems are beyond the reach of all but a handful of users. Almost all depend upon some type of video equipment. This minimally requires a camera and frame-grab capability in every office (as with PORTHOLES), to the extreme of full-bandwidth video, audio, and software-controlled video-switching (as with CRUISER). Similarly, active badge locators are available in only isolated laboratories, and it is unclear if they will attain widespread acceptance. While not condemning video and active badges as a communication mechanism, we believe that the base technology of "making contact" should cater to the *lowest* common denominator, as high-end approaches will place unnecessary restrictions on whom one can get in touch with (violating most of points 1a-e in Table 1).

A second problem with some of these systems is that they are not well integrated into the software environment. Its users, even after initial contact is made, may still have considerable difficulty initiating any groupware applications (making points 2a-e in Table 1 hard to achieve). Those that do (such as MESSENGER) do not promote media-switching.

Third, many systems will only work with small communities and will not scale up particularly well. Because the Internet is vast and expanding, users must have the capability of selectively probing different sites, groups, and individuals (point 1b).

A fourth problem is that many systems assume that they somehow know how other people can be reached in principle—office locations, phone numbers, email addresses, and so on. Yet this is not necessarily true with Internet contact. We may know that a person is on the net somewhere, but have no idea what their personal or site address is.

AN INTEGRATED COMMUNICATION AND SOCIAL PRESENCE ASSISTANT

We are working on a system called TELEFREEK that indicates social presence of others and allows its users to find and select appropriate communication channels, according to the factors listed in Table 1. Because our current implementation is work in progress, it is crude and

Petitioners
Exh-1015, p. 006

does not yet support all the features we want. This section will therefore stress our design ideas, as we believe it is important to show where we are going and why we are going there. Later, we will provide a snapshot of what we have implemented to date.

Design Overview

The TELEFREEK design is founded on four primary observations.

1. The issues of social presence and integration of communication methods are closely related and mutually dependent. They are too tightly intertwined to be supported separately.
2. Previous systems providing social presence information (such as media spaces) have done so in too passive a manner. Their failure to explicitly cater for users' active social presence requirements is reflected by unexpected and inefficient use of communication media, as observed by full video being used only to see when someone returns to an office [12].
3. Most social presence and integration systems assume the availability of high-bandwidth communication mechanisms and specialized equipment, and have based their functionality around them [4]. They appear to surmise that media spaces are necessary to find partners and to supply a communication link. Little attention has been paid to the integration of existing low-bandwidth communication facilities, which is all that most people have available.
4. Many resources for communication, collaboration, and social presence are already freely available on the network. These offer an extremely valuable resource, but their lack of integration renders them inaccessible to the majority of computer users. Access to such facilities should be merged, and doing so promotes synergy in communication and collaboration support. It is natural that, for example, information *about* communications (email addresses, surface mail addresses, and telephone numbers) should be combined with *access to* communication mechanisms.

We want TELEFREEK to provide several fundamental capabilities, based as much as possible on facilities and information freely available on the Internet. These are listed below, and are referenced back to the points in Table 1.

- Users are kept aware of who is around at *particular* sites in the networked community (point 1a). Given the size of the Internet, awareness of the complete community is both impossible and undesirable.
- Users may apply custom filters to restrict their views of the community to particular people or to people with particular attributes (point 1b,c).
- Links to people who have recently communicated with the user are automatically maintained. This is a reasonable system heuristic for tracking people that the user is likely to want to contact (point 1c).
- Users can request detailed information about any person on the network (points 1b-e).
- Users can be notified of status changes about particular people, e.g. they can set up "ambushes" that wait for someone to log on the network (point 1c).
- Similarly, users can be notified of status changes about information and system state, e.g. incoming mail, news articles, and new groupware meetings. This idea recognizes that communication may be initiated by someone else and that we may be interested in that communication (point 1e).
- Users can add hooks to new communication media as they become available (points 2a-e).
- When a user selects a person, the system determines what communication channels to that person are reasonable, and makes them easy to establish (point 2a-e).
- The system tracks all the people it has seen, and automatically maintains up to date information about them. For example, every incoming email message contains a name and an email address, and sometimes physical addresses and telephone numbers. These are added to the address data base. Information can be further extended and updated by having systems exchange their databases with each other during the normal course of human communication (the *liveware* concept [35]).

We want TELEFREEK to draw on resources *freely* available to networked computers users—because there are no special requirements in setting it up, one user's rejection of TELEFREEK does not impact on the benefits available to others [8]*. It provides a heterogeneous portfolio of information about others, access to communication resources, together with social presence guidance and other facilities that make up the communication requirements detailed earlier.

Unlike the majority of CSCW applications, TELEFREEK does not directly provide a medium for collaboration. Rather, it provides access to a variety of media, and information guiding the selection of appropriate ones. The goal is to ease the individual's selection of whom to collaborate with, how to maintain the communication, and how to ease people's transitions between interaction mechanisms. It also provides an extensible interaction platform, where new media and resources can be added by the end-user.

A snapshot of the current implementation

There are three versions of TELEFREEK: a HYPERCARD mockup used to elicit user feedback, a second version that operates over a local area net community, and a third version (in progress) that works over the entire Internet. This paper will concentrate mostly on version 2.

The TELEFREEK-2 prototype runs under Unix and the X window system and uses common facilities available on the Internet (see Cockburn [7] for full details). To maximize

---

* Of course, this can be pushed only so far. While a good many capabilities are achievable "for free," others do require a special software installation at remote sites.

Petitioners
Exh-1015, p. 007

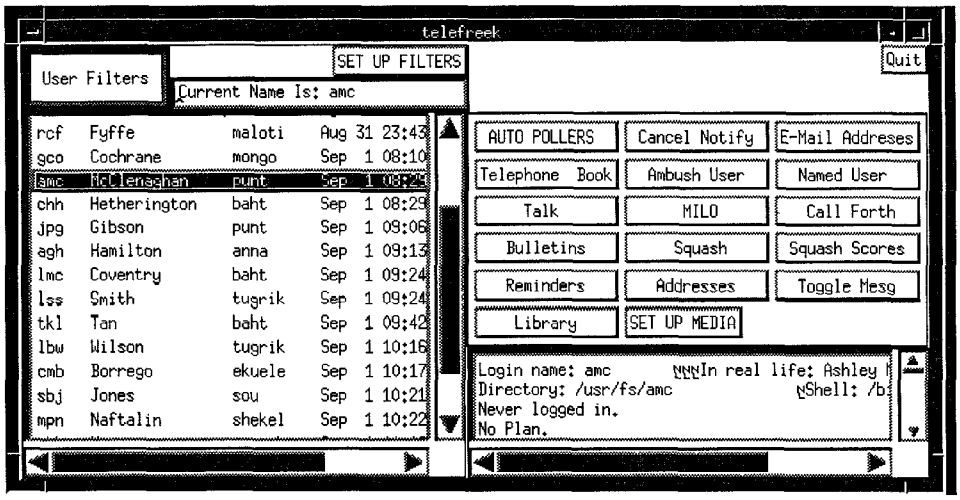

Figure 1. The main TELEFREEK display, showing user `McClenaghan` selected for inquiry and interaction.

portability and flexibility, almost all processing of TELEFREEK's utilities is delegated to sub-units independent of the core system (in the Unix environment these units are shell scripts).

Figure 1 shows the main window of TELEFREEK. The left hand side deals primarily with social presence information, while the right offers access to a variety of heterogeneous communication and subsidiary information facilities. The four major interface components of TELEFREEK are described below.

The Community List. The scrollable list on the bottom left shows a filtered view of the user community (filters will be discussed later). Each line, which is updated every few minutes, contains the login name of the person, their real name, the machine they are on, their login time, and (just off the view) how long they have been idle. When an individual is selected, any extra information known about the user is displayed in the bottom right window (by default, the system will automatically 'finger' that user and show the results). TELEFREEK will also set that user as the current candidate for communication.

User Filters. The button on the top left of Figure 1 pops up a menu of items describing different community subgroups, such as Friends, All Local Users, C Experts, CSCW Group, and so on. Selecting an item updates the community list to show only the users in that subgroup.

What makes this a powerful facility is that the set of filters available throughout this menu can be altered and extended by the end user though the Set Up Filters button (top center of Figure 1). This pops up a form that ties a menu item with a variant of the Unix *rwho* command that accepts arbitrary user names as arguments. The information returned by the command is then massaged and placed into the community list. While this means that a user must specify all members in each subgroup, future versions of TELEFREEK will generalize this capability to search for user attributes.

Communication and information facilities. The array of buttons on the right of Figure 1 provides access to most of TELEFREEK's communication and information facilities. Buttons generally fall into several categories or have specific functions.

- *Communication mechanisms.* Buttons can establish contact across a particular media channel with the user(s) currently selected in the community list. Selecting Talk, for example, will establish a two-way real-time text dialogue with the selected user. Similar functions could be added for email, and, given the equipment, video-call facilities.
- *Groupware applications.* Buttons establish task-specific groupware applications between users. For instance the Milo button launches the MILO co-authoring system [21].
- *Establishing links* to external sites. The buttons Call Forth and Library execute the command sequences to dial up external accounts and display the communication in a pop up window.
- *Accessing communication information* resources. Telephone Book and Email Addresses look up the current user in a data base, and show the results in the information display.
- *General communication and social functions.* The buttons Reminders, Bulletins, and Squash launch or display Unix applications in pop up windows, or display their output in the general TELEFREEK information display. The point is that access to these general social awareness utilities (for instance encouraging social interaction through the squash challenge ladder) is integrated within TELEFREEK; users will be more likely to keep up to date with such social activities if the overheads of doing so are reduced.
- *Interrupt status.* Toggle Mesg toggles the ability of others make contact through the TALK facility. A message reporting the current access status is provided in the information display.

Other specific facilities provided by TELEFREEK follow.

- Ambush User provides notification when an individual

Petitioners
Exh-1015, p. 008

or any member of particular sub-groups arrives. This is a base-technology variant of the ambush phenomenon observed in CRUISER use.

- Auto Pollers access user customised facilities that periodically check for status changes, such as the arrival of mail or a bulletin announcement. Announcements are made by an icon change, by an audible beep, and by printing the relevant event in the information display.

Extensibility and customization. As with filters, the power of buttons is that they can be defined and customized to fit the particular needs of the user. While some of the basic buttons are installed by system administrators, others can be hand-crafted. A user can select the Set Up Media button, and can add an entry to a form that specifies the button name, the Unix facility it should invoke, how it integrates with information known to TELEFREEK, and what TELEFREEK should do with the results. We believe this critical because the facilities available on a network are not static; new ones become available, and colleagues find more efficient ways of doing things. For example, digital audio is just now becoming available on many workstations, and network-based "telephone calls" are possible. It would be a fairly simple matter to add a new button to TELEFREEK that invokes the digital phone program with the currently selected user.

We also expect that people will actively share the facilities (the buttons and filters) that they find particularly useful [27]. Since all information is stored in TELEFREEK as human-readable text, extensions could be passed on through email. When facilities become generic, they could be installed within the TELEFREEK core as defaults by system administrators.

Integration. Not yet implemented in TELEFREEK but tested in HYPERCARD mockups is a way to integrate incoming communication, and a way to show users what communication channels are appropriate. For example, say a user was notified by TELEFREEK of new mail. When the mail is read, TELEFREEK will automatically display all it knows about the sender, and make that person the current candidate for communication (i.e. as if that person was selected from the community list). The person's name, email address, and other essential information are then automatically added to the appropriate data bases (such as those contained by the Telephone Book and the E-Mail Addresses). The system then sees if that person is logged on. If they are not, then the buttons invoking real-time computer dialogs are disabled. If they are on the system, those buttons are enabled. Similarly, if the system knows the telephone number of that person, a button would be enabled that could allow one to place the call.

CONCLUSIONS

This paper argued that there are several important factors for getting a group communicating through groupware, the most basic being how to get in touch with the right people, and how to choose an appropriate communication system. We believe the factors discussed under these headings form the minimum design foundation for any infrastructure supporting groupware on a wide area network.

While several CSCW researchers have tackled these issues, they either treat them in isolation or propose an approach requiring a very high technological entry level that is beyond the reach of most computer users. Under the guise of a design summary, we presented several specific capabilities—based on the human factors issues identified previously—of a system that integrates communication and social presence. A premise of the design is that its core would use as much as possible facilities and information freely available on Internet. Because a full blown contact facilitation system would require some information that is not available "for free," it should be customizable to take into account features with restricted availability or limited uses.

The paper also presented our work towards TELEFREEK, a flexible, extensible, and customizable platform that makes it easy for people to contact others, and that integrates people's access to a variety of communication facilities. It shows that such systems can be constructed without resorting to high-end technology. While not as glamorous as media spaces and active badges, it gets the job done. Of course, TELEFREEK is still in progress, and work remains to smooth out its interface and to meet all our design requirements.

The paper did not discuss privacy violations. There is a very fine line between contact facilitation (which could help people) and surveillance (which intrudes on people's privacy). While simple ideas such as reciprocity [12] and user control of their environment [25] can alleviate a few concerns, it is a difficult subject that must be addressed from technical, social, organizational, and legal perspectives [6].

ACKNOWLEDGMENTS

The Natural Sciences and Engineering Research Council of Canada and the Isle of Man Education Department provided partial funding for this project. Ernie Chang, of the Alberta Research Council, motivated us with his insights into the problems of making contact. Sean Brandenburg kept us up to date with his survey of awareness systems.

REFERENCES

1. Ackerman, M. (1991) "Shared expertise and the Answer Garden." In *Proceedings of ACM CHI'91 Conference on Human Factors in Computing Systems*, pp. 489, New Orleans, Louisiana, April 28-May 2, ACM Press.
2. Baecker, R. (1993) *Readings in groupware and computer-supported cooperative work: Assisting human-human collaboration*, Morgan-Kaufmann, San Mateo, CA.
3. Bair, J. H. (1989) "Supporting cooperative work with computers: Addressing meeting mania." In *Proceedings of the 34th IEEE Computer Society International*

Petitioners
Exh-1015, p. 009

Conference—CompCon Spring, pp. 208-217, San Francisco, CA, February 27-March 3.

4. Bly, S. A., Harrison, S. R. and Irwin, S. (1993) "Media spaces: Video, audio, and computing." *Comm ACM*, **36**(1), pp. 28-47, January.
5. Buxton, W. (1992) "Telepresence: integrating shared task and personal spaces." In *Proceedings of Graphics Interface*, Vancouver, Morgan-Kaufmann.
6. Clement, A. (ed). (1993) "Special issue: Workshop and panel of privacy considerations, CSCW'92." *SIGOIS Bulletin*, **14**(1), pp. 3-42, July.
7. Cockburn, A. J. G. (1993) "Groupware design: Principles, prototypes and systems." PhD thesis, Department of Computer Science, University of Stirling, Stirling, Scotland, April.
8. Cockburn, A. J. G. and Thimbleby, H. (1991) "A reflexive perspective of CSCW." *ACM SIGCHI Bulletin*, **23**(3), pp. 63-68, July.
9. Dourish, P. and Bly, S. (1992) "Portholes: Supporting awareness in a distributed work group." In *Proceedings of the CHI'92 Conference on Human Factors in Computing Systems*, pp. 541-547, Monterey, May 3-7, ACM Press.
10. Ellis, C. A., Gibbs, S. J. and Rein, G. L. (1991) "Groupware: Some issues and experiences." *Comm ACM*, **34**(1), pp. 38-58, January.
11. Fish, R. S., Kraut, R. E. and Chalfonte, B.L. (1990) "The VideoWindow system in informal communications." In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '90)*, pp. 1-11, Los Angeles, California, October, ACM Press.
12. Fish, R. S., Kraut, R. E. and Root, R. W. (1992) "Evaluating video as a technology for informal communication." In *ACM/SIGCHI Conference on Human Factors in Computing Systems*, Monteray, California, May, ACM Press.
13. Gale, S. (1989) "Adding audio and video to an office environment." In *Proceedings of the 1st European Conference on Computer Supported Cooperative Work (EC-CSCW '89)*, Gatwick, U.K., September 13-15, Computer Sciences House, Sloug, UK.
14. Gaver, W., Moran, T., MacLean, A., Lovstrand, L., Dourish, P., Carter, K. and Buxton, B. (1992) "Realizing a video environment: Europarc's rave system." In *Proceedings of ACM CHI'92 Conference on Human Factors in Computing Systems*, pp. 27-35, Monteray, California, May 3-7, ACM Press.
15. Goodman, G. and Abel, M. (1987) "Communication and collaboration: Facilitating cooperative work through communciation." *Office: Technology and People*, **3**(2), pp. 129-146.
16. Greenberg, S., Bohnet, R., Roseman, M. and Webster, D. (1992) "GroupSketch." *SIGGRAPH Video Review (supplement of the ACM Transactions on Graphical Systems)*, **87**, November. A video in the ACM CSCW '92 technical video stream.
17. Greenberg, S. and Chang, E. (1990) "Computer support for real time collaborative work." *Congressus Numerantium*, **75**, pp. 247-262, June.
18. Harper, R. H. R., Lamming, M. G. and Newman, W. M. (1992) "Locating systems at work: Implications for the development of active badge applications." *Interacting with Computers*, **4**(3), pp. 343-363, December.
19. Hopper, A. (1993) "The Active Badge system." *SIGGRAPH Video Review (supplement of the ACM Transactions on Graphical Systems)*, April. A video in the ACM INTERCHI'93 technical video stream.
20. Johansen, R. (1988) *Groupware: Computer Support for Business Teams*, The Free Press, Macmillan Inc., NY.
21. Jones, S. (1993) "MILO: A computer based tool for (co)authoring structured documents." In *Computer Supported Collaborative Writing*, M. Sharples ed. Springer-Verlag.
22. Kraut, R., Egido, C. and Galegher, J. (1988) "Patterns of contact and communication in scientific collaboration." In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '88)*, pp. 1-12, Portland, Oregon, September 26-28, ACM Press.
23. Kraut, R. E. and Streeter, L. A. (1990) "Satisfying the need to know: Interpersonal information access." In *Proceedings of Interact '90*, pp. 909-915, Elsevier Science Publishers B.B. (North Holland).
24. Li, J. and Mantei, M. (1992) "Working together, virtually." In *Proceedings of Graphics Interface '92*, pp. 115-122, Vancouver, Morgan-Kaufman.
25. Louie, G., Mantei, M. and Sellen, A. (In press) "Making contact in a multi-media environment." *Behaviour and Information Technology*.
26. Lovstrand, L. (1991) "Being selectively aware with the Khronika system." In *Proceedings of the European Conference of Computer Supported Cooperative Work (ECSCW '91)*, pp. 265-277, Amsterdam, Sept 25-27, Kluwar Press.
27. MacLean, A., Carter, K., Lovstrand, L. and Moran, T. (1990) "User-tailorable systems: pressing the issues with buttons." In *Proceedings of ACM CHI'90 Conference on Human Factors in Computing Systems*, pp. 175-182, Seattle, Washington, April 1-5, ACM Press.
28. Mantei, M., Baecker, R. M., Sellen, A. J., Buxton, W. A. S. and Milligan, T. (1991) "Experiences in the use of a media space." In *ACM SIGCHI Conference on Human Factors in Computing Systems*, pp. 203-208, New Orleans, April 28-May 2, ACM Press.
29. Newman, W. N., Eldridge, M. and Lamming, M. G. (1991) "Pepys: Generating autobiographies by automatic tracking." In *Proceedings of the European Conference of Computer Supported Cooperative Work (ECSCW '91)*, Amsterdam, Sept 25-27, Kluwar Press.
30. Prinz, W. and Pennelli, P. (1989) "Relevance of the

Petitioners
Exh-1015, p. 010

X.500 directory to CSCW applications: Directory support for computer based group communication." In *Proceedings of the 1st European Conference on Computer Supported Cooperative Work (EC-CSCW '89)*, Gatwick, U.K., September 13-15.

31. Root, W. R. (1988) "Design of a multi-media vehicle for social browsing." In *Proceedings of the Conference on Computer-Supported Cooperative Work (CSCW '88)*, pp. 25-38, Portland, Oregon, September 26-28, ACM Press.
32. Stults, R. (1988) "Experimental uses of video to support design activities." Research report, Xerox PARC, Palo Alto, California.
33. Tatar, D. G., Foster, G. and Bobrow, D. G. (1991) "Design for conversation: Lessons from Cognoter." *Int J Man Machine Studies*, **34**(2), pp. 185-210, February.
34. Want, R. (1992) "The active badge locator system." *ACM Transactions on Information Systems*, January.
35. Witten, I. H., Thimbleby, H. W., Coulouris, G. and Greenberg, S. (1991) "Liveware: A new approach to sharing data in social networks." *Int J Man Machine Studies*, **34**(3), pp. 337-348, March.

Petitioners
Exh-1015, p. 011