"Improving Communications at the Desktop", A. Hopper, Report of the Royal Society Discussion Meeting, "Communications After 2000 AD" held on 18/19 March 1992. Published by Chapman and Hall, 1993.

16

Improving communications at the desktop

Dr A. Hopper, University of Cambridge, and Olivetti Research Ltd

16.1 MULTI-MEDIA COMMUNICATION SYSTEMS

On a typical desktop of today we find a telephone and a computer. A regular telephone is used for audio and the computer provides communications in the form of electronic mail and sometimes facsimile. We may think of the user interface as being an electronic desktop. In future, the telephone will provide additional features and higher quality sound. It will probably still be quite small and with some form of display. The workstation will also be able to handle audio, but additionally video will pass through it. Such multi-media devices are likely to be well connected to others, both in the local area and the rest of the planet. Figure 16.1 shows some of the networks, devices and applications we may expect to find in communication systems which combine various form of media, in particular audio and video (I.E.E.E., 1990; 1991*a*).

Networks

Because audio and video applications are considered important, many networks have properties suitable for multi-media designed in from the outset. It is advantageous if the network provides good jitter control so that synchronization can be moved to the source end. In general, interleaving at the frame level is not adequate with normal frame rates and interleaving has to be done at the sub-frame or tile level. The asynchronous transfer mode (ATM) approach is well suited for this purpose and may become widely used. ATM networks of various kinds are being designed from infra-red systems to electronic systems using high capacity switches.

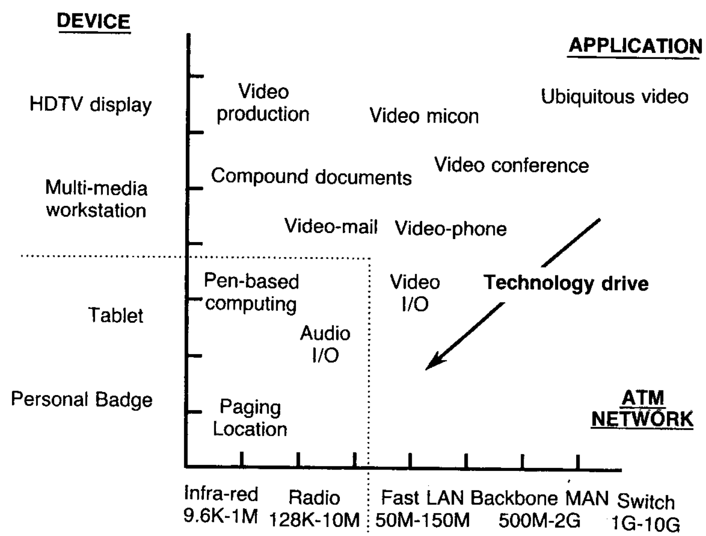

Figure 16.1 The multi-media communication space.

Devices

Devices carried by users are very much affected by their physical size and convenience of use. The smallest may be a personal badge which, used in combination with infra-red or radio systems, makes possible location and paging applications. A display tablet is larger and allows the pen-based model of computing to be used so that communication combining audio with writing on a shared screen may be possible. Until recently, workstations required special hardware to handle audio and video. A threshold is now being crossed where both CPU power and storage capabilities on standard machines are such that vision and hearing can be satisfied using largely non-specialized, and thus volume-produced, hardware. A workstation attached to a fast LAN can provide video-mail and video-phone applications, while a more complex workstation will be able to handle many multiple streams in parallel. A combination of a very high speed network and a HDTV workstation display will enable multiple video streams to be handled in a totally omnipresent way.

Storage

The storage devices provided for multi-media systems will be used both for storage of conventional data and for handling of synchronized streams

for real-time applications. The design of such storage systems has to be oriented both towards bursty high speed transfers to one device at a time, and also to well controlled parallel delivery of continuous streams. The digital representation of multi-media signals will mean that files have to be arranged in ways which gives fast random access. The spare CPU power of such file servers will be used to improve the quality of the stored images by software digital signal processing. It may also be used for generation of navigation information and annotation of real-time streams to make subsequent retrieval easier.

Compression

Minimizing the information loss at the transmitter allows the receiver better control of how to interpret, scale or use the data. In some applications, such as facsimile, only point-to-point transmission is envisaged and so a large amount of compression can be achieved for that single transfer. In computer systems arbitrary transfer of data between different points takes place and thus it is crucial that the compression systems should make this possible. In particular, it is necessary to consider the case where a particular stream is passed several times through a cycle of compression, decompression, and manipulation. Information loss has to be minimized or made to approach some asymptote. Compression can also make it difficult to combine different types of data such as video and graphics.

Types of traffic

Because network traffic needs to be real-time only if an individual is viewing it, other transfers can take place at either higher or lower than real-time speeds. If the camera becomes widely used the amount of video and audio data that will be present will greatly increase. This bulk data will be transmitted across networks in conventional ways and may dominate. Another type of traffic will be location information used to indicate the presence or movement of objects. This traffic has its own real-time restrictions because delay in reporting movement of objects may render the information useless.

16.2 DESKTOP CAMERA

A project started in 1987 at the Olivetti Research Laboratory and the University of Cambridge Computer Laboratory has dealt with the design and use of systems for desktop multi-media applications. The objectives of the Pandora project were to investigate how to construct such systems and also to find out if using a camera or microphone on the desktop as part of the normal computer environment is useful.

16.2.1 Pandora architecture

The Pandora Project set out with a number of goals and constraints. It was decided that the system would be completely digital and that the video and audio would be of medium quality. No hard restrictions were to be placed on the number of real-time streams it could handle. All streams were to be networked and switched in unconstrained ways subject to the limit imposed by total system performance. The system was based on a Unix workstation and so the screen for delivery of images was the same as one for the delivery of windows and conventional computer applications (Hopper, 1990; 1991).

In order to implement a flexible system it was necessary to adopt a hardware-oriented approach. This involved the use of a video and audio subsystem called Pandora's Box attached to the Unix workstation. For each pixel Pandora's Box can choose to pass through the original workstation pixel or to introduce a pixel of its own. There is no restriction on the pattern and shape of replaced pixels although in practice only rectangular video windows are used. As well as handling input from a camera, Pandora's Box interfaces the microphone, speaker and network. The hardware itself consists of several transputers, controlling capture, display, audio and network cards. A low level kernel provides the switching control. The kernel has a general purpose interface through which the workstation can request streams to be set up as required. An ATM networking system was adopted with a capacity of 500 Mbps in the backbone with distribution on local sites being done at 50 Mbps (Greaves *et al.*, 1990).

Table 16.1 shows the bandwidth requirements of the more commonly used video and audio formats on the Pandora system. When the streams are transmitted, an overhead of about 5% has to be added for protocol headers. The quality of the images was chosen so that the ATM network could support many such streams without interference. Similarly the workstation display can support many streams without constraint. The bottleneck in the system is the interface between Pandora's Box and the network; this can support about 5 Mbps of real-time traffic in arbitrary combinations. As well as handling video, Pandora's Box provides audio mixing so that up to about five digital audio streams can be received by the box and mixed correctly for output through a single loudspeaker.

Figure 16.2 shows the Pandora System in use at present. About 25 Pandora's Box workstations are dispersed across two sites, half a mile apart. The Pandora System provides services such as TV, multi-media file servers, and bridges to conventional LANs and wide area networks. The system has been in use for about 18 months in a community of about 50 users of which half use it regularly.

Table 16.1 Pandora stream properties

	size	*frame rate*	*compression factor*	*required bandwidth*
normal video	256 × 256 × 8	25	8	1.64 Mbps
medium video	128 × 128 × 8	25	4	819 Kbps
small video	64 × 64 × 8	12.5	2	205 Kbps
audio	8 bit μ-law	8K		64 Kbps

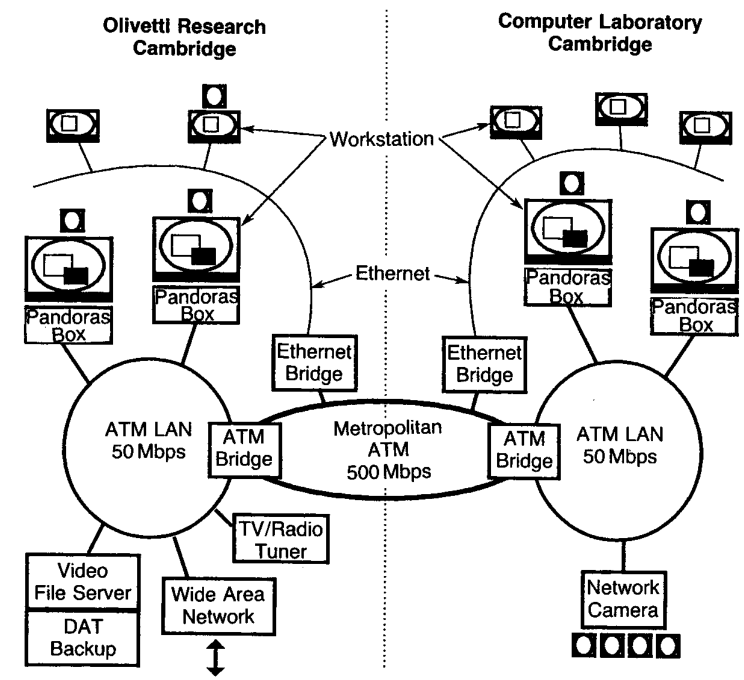

Figure 16.2 Pandora system.

Table 16.2 Summary of applications

General points		*Acceptability*
local image not mirrored		high
lip-sync with own image		?
no eye contact		medium
low quality threshold		high
Applications		*Popularity*
Using live streams:	2-way video-phone	high
	3-way, 4-way video-phone	low
	live information services (TV, radio)	low
	permanent video links between sites	medium
	peek in office (bidirectional)	medium
	look everywhere	low
Using storage:	video-mail	very high
	system documentation	low
	general material	medium
	personal introductions	low
	latest news	medium
	mixed text/video mail	very low

16.2.2 Pandora applications

The way the Pandora system and applications have been used is summarized in Table 16.2. With a small community of users it is not possible to extrapolate to larger groups, although it is likely that some of the experience will be repeated elsewhere. However, many of the predictions we made have, in practice, turned out to be inaccurate and the new facilities became popular for quite unexpected reasons.

General

At first users may have to rearrange their office to give good lighting and they will have to become accustomed to a camera. Pandora applications seem more natural in hands-free mode which requires careful consideration of the acoustic properties of an office or a room. The user speaks towards the workstation which has a microphone and a loudspeaker nearby. Feedback control and echo-cancelling circuitry is used to improve sound quality.

In the Pandora System a local image is shown as a true image and not a mirror image. A user of the Pandora System quickly gets used to not

treating the screen as a mirror. This is fortunate because a true image has to be transmitted and generating both true and mirrored images in hardware would be expensive. Inevitably there is a slight delay between movement of the lips and the electronic version as presented on the local screen. Networked streams showing other parties can be well synchronized.

Because conventional workstations are used, the cameras are to one side or above the screen rather than being behind it. The resolution is such that lack of eye contact is not immediately apparent. Users are prepared to accept a lower quality threshold than was anticipated. High frame rates with lower resolutions are preferred to better quality images which appear jerky because of low frame rates. Break up of the audio streams renders the applications unusable.

Applications using live streams

The video-phone is the most popular application using live streams. In a video-phone the user sees a local image of himself and uses this to centre his head on the screen. Most users prefer to see a small image of themselves, finding a large one distracting. A large image of the other party is shown in another window and hands-free operation is almost always used. The receiver is invited to join a video-phone conversation by a window appearing on the screen indicating that a call is being requested. There is a choice of several replies: accepting the call, asking for the call to be made again later, or indicating that a return call will be made. This has turned out to be a particularly useful feature and makes a video-phone a more acceptable piece of office equipment.

Two way video-phone conversations are popular and typically last longer than telephone conversations for the same users. At Pandora resolutions the body language is passed over well, and it is possible to tell whether the correspondent is still interested in what is being said. In effect, conversations can last for quite a long time because of this feedback. Figure 16.3 shows the distribution of video-call lengths. Most calls last between 1 and 5 minutes although there is quite a long tail in the distribution and some calls go on for 15 minutes or more. The shape of this graph has largely remained unchanged with use of the system. Three and four way video-calls are less popular, but this may be because the community of users is not large enough.

Another live application is the provision of media information services which combine video and audio. Initially it appeared that these would be popular but in practice they are hardly ever used except in times of crisis or some significant news event.

It is possible to set up a video link between different places and leave it

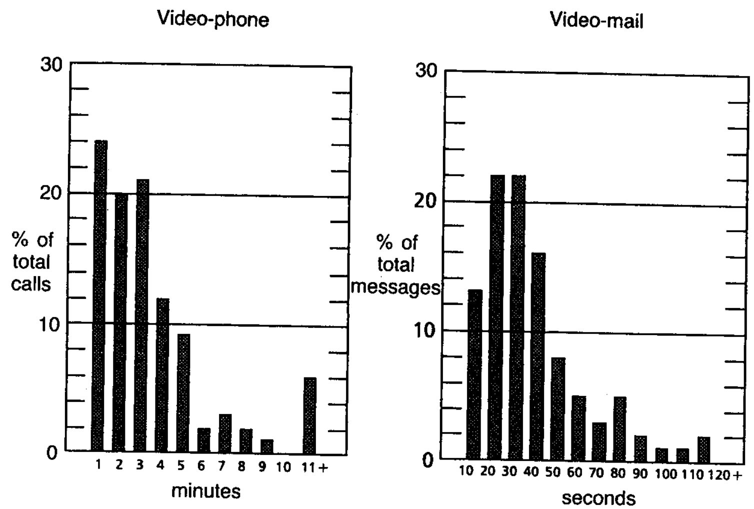

Figure 16.3 Pandora statistics.

running. Normally these are bi-directional and it has been observed that from time to time such links are set up between public areas. The users comment that even though they know where everybody is (see Active Badge section later) it is nevertheless useful to have some feeling for what is going on in another room. When invoked between private offices bi-directional video links can be used to peek into another office. The parties see a picture of each other on the screen. This is usually followed by a smile, a wave and shutdown after 3 or 4 seconds and is moderately popular.

It is possible to receive video streams from all cameras in the system at once. This application is only used for demonstrations and does not seem to fulfil any useful purpose.

Applications using storage

The multi-media file server provides a general purpose storage facility for audio and video. It can support about five simultaneous streams on the network interface. The most popular application using the file server is video-mail, which involves the recording and transmission of a video/audio message. A delivery message is sent to the receiver and the video-mail is viewed when convenient. Most messages begin with a courtesy and then have a substantial part with a nuance or several conditional points

which, if written, would require careful drafting. Figure 16.3 shows the distribution of video-mail message lengths. As with video-phones the shape of this graph has not changed substantially in the lifetime of the system. Video-mail longer than about two minutes is considered unduly long although there is no upper limit imposed by the system. It is interesting to note that no video-mail 'flaming' takes place. For a number of users the video-mail system has become the main communications channel although they use text where accuracy is required. Video-mail is quite different to and complements other forms of communication.

The file server currently has 2.5 GBytes of storage which represents about 6.5 hours of typical recordings. There are no restrictions placed on how this facility is used and more storage could be made available if required. About 20% of the contents represent video-mail which has a short lifetime of a few days for each piece. Material which is available to everybody and kept much longer includes systems documentation and general items. Some users like to record personal introductions which contain any message about themselves they wish. Thus while the great majority of recordings that have been made represent video-mail the storage system is mostly filled with other material.

There are a number of applications which combine live streams and storage. One of these is the automatic recording and re-recording of news, weather and other live programmes to the file server. The file containing latest news is updated at regular intervals throughout the day. Thus every user has the latest news available at all times. It is the ability for the user to choose the time when to view that makes this application more popular than the live service.

It has been a surprise that a system which allows users to combine audio, video and text has not proved popular. In this application the normal mail handler is used to generate text into which video clips are inserted as required. The whole compound document is sent to another party who can read the text and view the video parts. There may be a number of reasons why this application is not popular. The first is that the manipulation tools which are presently available are not good enough to make it easy to generate such documents. An alternative view is that video-mail and text-mail provide different methods of communication and it is quite natural to make a choice of the one better suited at a particular time. For the majority of communications the video-mail system is appropriate because it is quicker, but where particular accuracy is required text-mail is used instead.

The Pandora environment enables new multi-media communication applications to be evaluated. For example an evaluation is being made of switching of additional streams during video-phones so that video-mail or other material can be shown in a video-phone session. However, at present re-use of material on the file sever is difficult because there is a lack of retrieval and manipulation techniques for multi-media data.

16.2.3 Extending the use of desktop cameras

Because video-mail does not require a high speed communications path from source to destination it is possible to deploy a system which uses the existing networks for distribution. The recording is made to a local disc (or network server) and from there the data is sent as a normal file transfer. The video-mail can be viewed on standard computer workstations because many have audio output and a fast enough frame store. To generate video-mail on standard workstations a simple capture card for digitizing the camera and microphone output is required. At present experiments are taking place with use of video-mail between research laboratories in Cambridge, (England) and the USA. It is too early to say how this use of video-mail will evolve but early signs are that it will be popular.

When the Pandora system was being designed, the multi-media features were only feasible for fixed workstations. However, it is now becoming possible to design a portable system of similar performance. This would extend use of Pandora applications in the local area. It will be interesting to see what types of communication device users are prepared to carry.

Recently CCD technology has become sufficiently mature to appear in combination with other logic on single chips. This will make it possible to design units with on-board digital processing to provide digital video in very flexible ways. Cameras will become ubiquitous, making feasible applications from simple observation systems to complex data gathering systems. A new project called Medusa aims to extend the use of cameras at the desktop and in a system. The assumption is made that many more cameras will be associated with each user and multiple networked streams will be generated. Thus a workstation might have eight or sixteen cameras, while a video system used in public areas might have many more.

In the Medusa system the receiver will have control over all streams and, subject to the sender's authorization, will be able to transmit them to a recipient for selection of which one to watch. The recipient may be aided by local hardware, software or hints from the transmitter in choosing the required stream. The principle of receiver control is important and applies not only to choices between multiple video and audio streams but also to representations for data of all types. While at present this is technically difficult because the data rates generated may be high, in due course, when the capacity of networks increases, it may not be a problem.

A multi-stream capability of this kind can be used in a number of ways. If the cameras are attached to a workstation directly then each one could point in a slightly different direction so that centring the head can be automatic. Alternatively, the cameras could be arranged to cover a particular volume of space so that the user can move around and think of the cameras as being omnipresent in that space. Another way of using multiple streams is to use each one for different types of information. For example, some of the cameras could be sensitive to the infra-red and may be able tell where hotter parts of the image are, which may be useful for applications such as tracking of faces or other objects. It may be possible that an infra-red grid can be projected and then viewed by a infra-red camera to obtain depth information from the distortions. Yet another way of arranging the cameras would be for use with three-dimensional displays. This would require all the sources to be set up in a coherent way such that the requirements of the 3-D display are met. A more conventional alternative for Medusa would be to use all the capacity to drive a single HDTV display.

16.3 DESKTOP INFORMATION

A way of improving communication at the desktop is to give a better idea of who is available, who is not, and what the user communication preferences and choices are. The Active Badge project aims to investigate the design and use of a small personal device to make more attractive the use of communications (and computer) facilities.

16.3.1 Active badge system

The Active Badge is used to provide information about where people are (Want *et al.*, 1992; Want and Hopper, 1992). It is battery powered, transmits in the infra-red spectrum and is approximately 60 × 60 × 8 millimeters. The transmissions take place every 15 seconds and identify the badge. There is a single button on the badge which causes an immediate transmission. Receivers are linked by wire to a computer and are placed so as to define cells for the coverage required. Normally they correspond to spaces occupied by one or a number of people. The range of the infra-red system is about 20 metres and communication does not need to be line of sight. The badge has a light dependent resistor used to reduce power consumption by decreasing the frequency of transmissions when in the dark. This also means that the user can switch the badge off by placing it in a pocket or face down on the table. Not all badge transmissions are picked up by a receiver, but by using simple algorithms in the receiving software the system can be made sufficiently accurate to be very useful

Table 16.3 Active badge display

Name	*Telephone*	*Position*	*Seen*	*Status*
P Ainsworth	343	Accounts	Static	Alone
M Chopping	410	R410 MC	Friday	
D Clarke	316	R316 DC	12:30	
D Garnett	218	R435 DG	12:20	
T Glauert	232	R310 TG	35 mins	
S Gotts	0	Reception	Static	With Jackson
D Greaves		Floor 3 Corridor	Moving	Alone
A Hopper	334639	Univ Comp Lab R76	Static	
A Jackson	0	Reception	Moving	With Gotts
A Jones	210	Meeting Room	Static	At a gathering
T King			Yesterday	Away in Italy
J Martin	310	Machine Room	24 Dec	
O Mason	210	Meeting Room	Static	At a gathering
D Milway	BUSY	R211 DM	Static	Alone
J Porter	398	Library	Static	Alone
C Turner		Front Door	Yesterday	
R Want	308	R215 RW	7 mins	
M Wilkes	210	Meeting Room	Static	At a gathering
S Wray	204	R212 SW	Static	Alone

14:09 Tuesday 8 January 1991

Many different applications which use the badge sighting information can be devised. One of the first was the provision of location information about individuals. A typical interface is shown in Table 16.3 and gives the name of the person, his position and the number of the nearest telephone. The column marked 'seen' makes it easy to distinguish between individuals who are static and those who are in the process of moving from one place to another. If the badge has not been sighted for three minutes, this column shows how many minutes ago the badge was last sighted (e.g., 35 mins), and after an hour changes to a time (e.g., 12:20). The 'status' column gives simple information about the users' (apparent) circumstances. The badge information is made available to all computer screens in the organization.

By pressing the badge button, a user can explicitly indicate that he does not wish to be disturbed (busy). This is cancelled automatically when the user moves away from the current cell. It is also possible for the user to specify a phrase to be displayed in the status column (e.g., Away in Italy). Thus when people go away they can indicate this to others. The message is cancelled automatically when they are sighted again. Using

automatic rules to cancel features is important because a principle of the design is to require the user to do very little to make it work. To generate location information wearing the badge is all that is necessary. With the 'away' feature an incorrect list of future absences does not matter.

Various versions of the user interface have been tried which include map displays, simple natural text generation systems and others. The Active Badge has also been used in a security role for example to control door locks and to blank a computer screen when the user walks away.

16.3.2 Use of the active badge system

The system has been in use for several years and continues to be very popular. Users interrogating the system have made it one of the biggest consumers of CPU cycles in an organization which uses computers for many applications. Figure 16.4 shows the proportion of time that people in the organization spend in their office and in the building. It can be seen that in this organization people spend a large proportion of their time away from their normal office and one of the reasons why the badges are successful is because less time is spent trying to find others. The most important single piece of information is that somebody is *not* available. A more subtle but equally important observation is that the badges provide a mechanism for people to be less intrusive with respect to others. Our initial thoughts were that the number of phone calls would go up because people could make contact more easily. In fact a self-filtering process

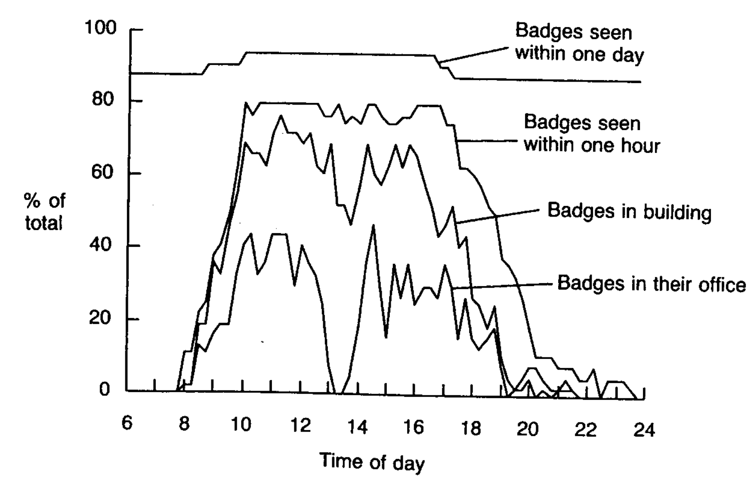

Figure 16.4 Mobility of badges.

takes place where people do not contact others if they see them in circumstances where they may not wish to be disturbed.

Having been first established on a site with 40 wearers (Olivetti Research), the system was extended to a second site with about 80 more wearers (University of Cambridge Computer Laboratory). Because the amount of data generated by movement of people is small, the speed of the communications channel required was not high. Within one organization, the wearing of a badge indicates the user is prepared to have information about their location distributed throughout that organization. Between the two organizations the opposite view is taken and it is necessary for each user to authorize explicitly their sightings being sent elsewhere. The two groups need to communicate frequently and in practice all those working on joint projects (and others) have authorized the information to flow in this way.

Other organizations which work collaboratively are now having the system installed (for example DEC Research). It will be interesting to see if inter-working across continents is also made easier using this approach.

The Active Badge location information has been integrated with the Pandora system. A video-phone can be routed to the Pandora's Box nearest to a user. Who is calling (not where they are) is indicated to the recipient which makes it easier to decide which response to give. It is also possible for both sides of a conversation to know who else is nearby to prevent confidential discussions taking place when other people are in ear shot but out of view. The badge data is also integrated with the video-mail system. Video-mail is most popular in a form where the user does not type anything and there is very little information to help subsequent retrieval. The badge data is used to annotate automatically the video stream. This permits the retrieval of all the video-mail files sent by a particular user, the ones that were recorded with visitors present, or any other category which can be derived from the badge annotations.

16.3.3 Extending desktop information

The cellular infra-red network can be extended to two-way communication which makes it possible for more sophisticated badges to be developed. However, the challenge is to design the simplest systems that provide useful features with only minimal or easy to learn interactions.

Two-way authenticated badges

A new badge has been developed which both receives and transmits. It is implemented using a normal microprocessor and its properties are controlled by software. The two-way badge has a tone generator and two LEDs making paging applications possible. The paging can be made location dependent so that for example a page as a reminder of a meeting only goes off if the user is not in the meeting room. The Authenticated Badge also has a 'yes' button and a 'no' button which the user can press to generate context dependent commands.

A transmit/receive active badge enables a more secure system to be devised by using a one-way authentication function. The badge identifies itself in the normal way but can also perform authentication by one-way mapping a 48-bit random number to another on demand. For more security it would be possible to incorporate a personal identification number which when revoked disables the badge.

Finally the Authenticated Badge has a simple radio receiver which enables a second type of location system to be implemented. A radio field can be associated with a piece of equipment. When a badge enters this field it is possible to determine within which one (if any) it is. The field can be made small so that when a user approaches a piece of equipment sufficiently close to use, this fact is made known to the system and the appropriate application invoked. The radio field is also used to trigger the badge when it is in low power mode between transmissions.

Badges for equipment

A badge suitable for attaching to equipment is being designed and has similar properties to the transmit-only Active Badge. It uses less power by only transmitting every five minutes and it can be attached in a way which indicates whether the equipment is switched on. By using the network of receivers, some objects can be located at the same level of granularity as personnel.

It is now possible to have an automatic inventory control system which keeps account of what there is and where it is. Tagging in this way will make possible an environment in which the location of all equipment, all I/O devices, all cameras and microphones, is known. With such an active environment it is possible to think of new applications which combine location information of personnel and of equipment.

16.4 THE ACTIVE OFFICE ENVIRONMENT

This paper has described how new communications facilities may be accepted by users. In a discussion entitled 'Technology in the Third Millennium' it is appropriate to speculate about future directions. The 'active office environment' is a framework for development based on experience with the Pandora and Active Badge projects.

16.4.1 Components making up the active office environment

The 'active office environment' will make widespread use of ATM networks. Fibre will be used both for long distance public systems and local private systems, while infra-red and radio will be used for portable applications. Wire drop cables will continue to be available for short distances because they will not have to operate at the same speeds as the underlying networks. Islands of high interconnection will be joined by high capacity links while elsewhere wide area radio and satellite systems will be used to provide some connectivity (I.E.E.E., 1991*b*; 1991*c*). A significant part of the traffic will be bulk image data which can be carried on conventional rather than ATM networks.

The devices using these communications system will be constrained by the facilities they provide and by size and convenience of use. They will all be capable of manipulating audio and video in the same way as any other type of data. A convergence of the way communications are handled within a device and on a network will take place so that there will be similar protocols and little mismatch for use of desktop or remote systems (Hayter and McAuley, 1991).

Automatic location of components will be used to offer users choices that are more comprehensive and easy. The Equipment Badge provides a particular level of granularity for doing this. However, other levels of location granularity may be available as well. Fine grain distribution of location information will need sophisticated algorithms for optimizing traffic flows. Systems which require location information infrequently will make enquiries on a need to know basis. Systems which require location information all the time, or cannot tolerate the delay associated with explicit enquiries, will require such information to be distributed automatically.

16.4.2 Personalization and mobility

A major difficulty when using computer and communications equipment is that user interfaces are not consistent. In the active office environment, the desktop will be made to follow the user and each piece of equipment will be personalized when used. One way this can be achieved is by using a badge. When a user approaches a piece of equipment he is associated with that equipment and the appropriate interface is invoked. To make this facility eventually possible on a world wide basis the Authenticated Badge transmits to the equipment a pointer (similar to an e-mail address) indicating the home base of the badge wearer. If the communications channel is of low bandwidth then the parameters which follow the user would only be of a parametric type. If the communications channel has

high bandwidth then all data can follow. An alternative and perhaps complementary way of providing the personalization facility would be to store the parameters describing the users preferences within the badge. An example of personalization in use would make the short codes on every telephone remain the same so that no matter where in the world the call was being made the same destination would be dialled.

Personalization can also be used to control the operation of the communications environment by incorporating the users wishes as specified in his user service profile. This will control distribution of location information, routing and forwarding of calls, and interactions with computer systems. Table 16.4 shows what might be a typical service profile. In this example the user has allowed location and status information to flow to one other organization and also to one individual no matter where that individual is. Call preferences are for the best quality available with the user choosing automatic forwarding and some automatic redirection. Simple interaction with the computer environment is allowed to indicate when video-mail arrives and when the user is in the wrong place with respect to an active calendar. Table 16.4 illustrates a text interface for making these choices as it may be that natural language systems operate well in this restricted domain.

Table 16.4 User service profile

Option	*Choice*
Distribution of location/status information	
Own organization	Yes
Other organizations	Computer Laboratory
Other individuals	No except Fred
Call control	
Video, Audio or Page	Best available
Fixed or portable	Fixed if nearby otherwise portable
Follow-me forwarding	Not if with the boss Not if in the library Not if in a meeting of 3 or more people Not if busy pressed on badge
Redirecting	To secretary in the morning Play video message 1 in the afternoon Play video message 2 if Fred
Computer Systems Interactions	
Page	Yes if video-mail arrives Yes if calendar exception

16.4.3 Communicating in the active office environment

A user will be able to identify the destination as a place, a person, or a facility. He will then choose the appropriate audio, video, or if real-time communication is not required, multi-media e-mail system. To do this, directory services will be provided which will deal with the recipients' mobility and preferences. The distribution of information will be subject to a hierarchy of access restrictions as specified in the user service profile (I.E.E.E., 1992).

There is a tension between the flexibility the receiver may wish to have by receiving all data and the security required by the source. If the caller has access to enough information about the recipient he may know when *not* to make the call and both parties may find this more appealing. When designing such systems it is a challenge to provide security features and methods of use which the users accept, so that information for making such decisions is generally available.

Directory services

Because video, audio and text will be available, one purpose of the directory services will be to provide information about the equipment at the destination and its compatibility with the source. A simple directory service of this type can be provided by relatively infrequently updated tables. A partially dynamic system will require frequent updates. This will enable the potential recipient to indicate status information such as away, portable not in use, or some other choice. A completely dynamic system would present the full picture to the caller including location of the recipient, of equipment nearby, and any other relevant information. This will require the directory to be completely dynamic, that is much of the data would have a short lifetime and have to be updated frequently. The operation of directories and the choices they give callers will be automatic. The implementation of the directory service would take into account the frequency at which entries change and how often they are used.

Call control

When the call is made it may be redirected as specified in the user service profile. If it is routed to the recipient there will be an indication of who is calling. A number of choices for dealing with the call will be available such as accepting the call, giving an immediate response to the caller (e.g., I will call you back), or using the redirection options indicated in the user service profile. If the call proceeds the specified media will be used.

If the recipient is not presented with the call or chooses to reject it an

automatic system for redirection will be used. The caller may be redirected in a number of ways or have a text, audio or video indication of what happened. In the example shown in Table 16.4 the recipient has specified that the secretary is to handle rerouted calls in the morning, a video message is to be played in the afternoon, and a special video message is to be played if a particular party calls.

16.4.4 Conclusion

As technologies develop the power of communicating devices will increase. They will also be able to communicate much better. The desktop will no longer be associated with a physical place nor an interface to a single system. It will become virtual and there will be one seamless interface which is consistent across devices and communication systems.

ACKNOWLEDGEMENT

My thanks go to the many colleagues both at Olivetti Research and the University of Cambridge Computer Laboratory whose joint efforts have made much of the work described in this paper possible. I am also grateful to Prof. Maurice Wilkes who made pertinent comments about this manuscript.

REFERENCES

Greaves, D.J., Lioupis, D. and Hopper, A. (1990) The Cambridge backbone ring. *Proc. I.E.E.E. Infocom'90*, San Francisco, June.

Hayter, M. and McAuley, D. (1991) The desk area network. *ACM Oper. Sys. Rev.*, October.

Hopper, A. (1990) Pandora – an experimental system for multi-media applications. *ACM Oper. Sys. Rev.*, April.

Hopper, A. (1991) Design and use of high-speed networks in multi-media applications, in *Proc. 3rd IFIP Conf. on High-Speed Networking*, Elsevier Science Publishers, March.

I.E.E.E. (1990) Multi-media communications, *I.E.E.E. J. Sel. Areas in Comm.*, April.

I.E.E.E. (1991*a*) The promise of the next decade, *Computer*, I.E.E.E. Comp. Soc., September.

I.E.E.E. (1991*b*) Satellite and terrestrial systems and services for travellers, *I.E.E.E. Comms. Mag.*, November.

I.E.E.E. (1991c) Architectures and protocols for integrated broadband switching, *I.E.E.E. J. Sel. Areas in Comm.*, December.

I.E.E.E. (1992) Intelligent networks, *I.E.E.E. Comms. Mag.*, February.

Want, R., Hopper, A., Falcao, V. and Gibbons, J. (1992) The Active Badge location system, *ACM Trans. on Inf. Sys.*, January.

Want, R. and Hopper, A. (1992) Active Badges and personal interactive computing objects, *I.E.E.E. Trans. on Consumer Elect.*, February.