The LAN as an Integrated Communications Environment

Roy Want and Andrew Hopper

Olivetti Research Ltd
Cambridge, England

For some time now it has been claimed that the Local Area Network (LAN) is suitable for the integration of voice and data media. Many projects have demonstrated the feasibility of these systems although only a few have become commercial LAN products.

The data transfer rates of LAN technology have been increasing in recent years to the point where 100 Mbs^{-1} is now realistic at a reasonable cost. The higher data rates allow the integration of video into the LAN environment without the need for complex compression techniques. Networks supporting small-packet transfers are particularly well suited to the real-time constraints of voice and video and can provide an environment that will support communication between multimedia workstations.

This paper describes a number of projects that highlight the advantages of using a LAN for the purposes of telephony and video transport, and as a means of distributing multimedia services around a site supplied by high-bandwidth ISDN links.

Roy Want was born in London, England. He studied computer science at Cambridge University and received a B.A. in 1983, and later a Ph.D. early in 1988. During this time his research centred around the Cambridge Distributed System and in particular a multimedia project called ISLAND (Integrated Service Local Area Network Development). He is now a research engineer at Olivetti Research Ltd in Cambridge and is involved in the research and development of multimedia workstations.

Andrew Hopper received the B.Sc. degree in computer technology from the University of Wales, University College of Swansea, in 1974 and the Ph.D. degree from the University of Cambridge in 1978.

From 1977 he worked at the University of Cambridge where he holds the position of Lecturer. His interests include local area networks, VLSI design, micro and multiprocessor design. He has many links with industry as a director of Acorn Computers Ltd., and Qudos Ltd., and from 1986 as Director of the Olivetti Research Laboratory, Cambridge, England. He is also a member of the Olivetti Research Board.

1 Introduction

The development of LANs with a 10Mbs^{-1} bandwidth in the late 70's brought with it the possibility of using one network to provide the transport medium for data and voice communication. Nowadays 100Mbs^{-1} is a practical speed for a LAN and this makes it feasible to extend this medium to video. Current research is considering networks in the gigabit range [Greaves 88] which will further enhance the usefulness of a LAN in the area of multimedia communications.

A single local network is desirable from the pragmatic perspective of wiring large numbers of devices together and providing the flexibility for reorganising their connectivity. One of the key properties of LANs suitable for multimedia research is that of low jitter and a guaranteed small upper-bound for delay in data transport. These are the properties of small-packet ATM networks and as a result they are suitable for the integration of voice and data traffic. The developments described in this paper have been carried out using *slotted rings* e.g. the Cambridge Ring (10Mbs^{-1}) [Wilkes 80], and later the Cambridge Fast Ring (100Mbs^{-1}) [Hopper 88], primarily for the reasons stated in the list below.

A slotted ring is an example of a small-packet ATM network and possesses the following useful properties:

- **Bandwidth is shared out equally between equal requestors:** this property results from stations passing on a slot after every transmission. The available bandwidth is also independent of the offered load. This is in direct contrast to CSMA/CD systems which experience a higher probability of packet collision with large loads and thus network utilisation decreases.

- **Low access delay:** slots are small packets of data that are passed between requestors, in the worst possible case, at the packet rate divided by the number of stations in the ring. In some Token Bus or CSMA/CD systems packets are considerably longer to reduce the overhead of the packet preamble. It is therefore, in general, not possible to guarantee the transport time for a packet and the requirements of voice cannot be met. In order to be able to handle voice these systems often have ad-hoc features that guarantee the transport delay of packets reserved for voice. However, in a small-packet ATM system this is a natural property of the network.

It can be concluded that if data users and voice users coexist on the same ring, each will have access to a slot with little delay. Moreover, the low-bandwidth voice user's requests will be met in full, leaving the remaining bandwidth divided among the needs of the data users.

2 Experimental Voice–Data Systems

There have been many projects that have demonstrated the feasibility of an integrated voice/data communication environment. One of the first research projects in this area was the Xerox PARC Etherphone project [Swinehart 83]. The principles of the Xerox architecture followed those of a distributed system inter-connected by an Ethernet. The three main components of the system were the Etherphone, the Telephone Control Server, and the Voice File Server. The Etherphone performed the voice digitising operation, and the transmission and reception of voice packets. The Telephone Control server responded to telephone key-strokes and was responsible for establishing connections between Etherphones. Because an Ethernet is not very suitable for voice traffic, these experiments were carried out on an isolated and unloaded 1.5Mbs^{-1} Ethernet. Additional control over the Etherphones from workstations was carried out through a gateway connected to the main distributed system. The project was a great success allowing sophisticated control of the Etherphone network from applications running in the Xerox development environment.

2.1 ISLAND

The ISLAND project [Ades 87] at the University of Cambridge was founded on similar principles to that of the Etherphone project, however with a 'slotted ring' a completely integrated voice-data network was possible. ISLAND set out to achieve three goals.

- Integration of voice and data traffic in one network.
- The development of a distributed PABX.
- Integration of voice and data services.

The first goal was satisfied by the use of a 10Mbs^{-1} Cambridge Ring as the communication network. The second and third required careful design. The following model was adopted for the PABX subsystem.

The telephones were of simple design and contained little knowledge of the overall system state. All services were designed to be as modular as possible. The ISLAND Exchange Server [Want 88] was replicated to ensure the high availability of a control service. Failure of support services could only cause a gradual degradation of the system service, but assuming the network was operational a telephone call could be made with high probability of success.

Considerable emphasis was placed on reliability because of the high expectations we have of existing telephone networks.

The ISLAND model tried to make as much use of existing network services as possible. To this end the File Server used by the Cambridge Distributed System was also used to store digital-voice recordings without any need for modification. A File Server is usually most efficient at storing and recovering large packets of data in a single transaction. Direct communication with a telephone implies handling small quantities of data periodically. By placing another service named a *Translator* [Calnan 87] between the two a compromise was found. The telephone could assume it was talking to another 'telephone like' service, but in reality the Translator was converting the voice samples to large packets to send to the File Server and, on reception of a packet from the File Server, breaking it into smaller packets to send to the telephone. If storage space was scarce and it was decided that the digital voice needed to be compressed, a compression server could be placed in the middle of the File Server–Translator data path.

Another example of a service provided by a PABX is the ability to set up conference calls. One way to provide this service in the environment of a system such as ISLAND is to use a broadcast protocol, however, this was not available for use in the Cambridge Ring.

It is possible to build a *Conference Server* capable of receiving voice packets from N servers, combining them and retransmitting the composite samples to all of the originating parties. Thus setting up a conference call from the point of view of the Exchange Server is only a matter of setting up N connections with the Conference Server, each connection using the same voice protocol as used between two telephones.

The approach used in the ISLAND project allows new services to be added to the network in a modular way, thus improving maintainability and providing a well structured framework to the system.

3 Distributing the PABX

The ISLAND project placed a great deal of emphasis on a simple telephone. ISLAND telephones could signal key-strokes and changes in handset state to a controller, and could display short text messages on a single line display or generate a variety of tones. In addition, each telephone had its own network interface, making it physically independent of the operation of any host computers that may be logically associated with it in an application.

A key part of the telephone operation was the use of a simple voice-protocol. Digital voice was sent as a stream of 2ms packets without error checking or retransmission. However, each packet contained a sequence number and a flag indicating silence or speech. The silence flag was used when a telephone was sending packets to a Translator in order to save disk storage space during silent periods. This meant 'silence' had to be reconstructed in the Translator on playback. A suitable level of digital white noise could be manufactured for this purpose. Error checking and retransmission is of little use in a voice protocol because of the real-time constraints on delay. Voice packets were chosen to be 2mS long because this size of voice segment can be lost on a random basis 1% of the time or less with a negligible loss of quality detected by the human ear [Gruber 83]. The sequence numbers allowed incoming packets to be placed at the correct position in a reception buffer in the event of lost data and additionally allowed the duration of silent periods to be calculated once active digital voice was stored on disk. Reasons for not adding complexity to a phone are:

- If too much functionality is placed in a telephone it fixes the operation of the overall system to the specification of the telephone.
- It increases its cost, complexity and need for maintenance. This factor maybe small for each telephone but there are a large number of telephones in a system.
- Simplicity gives a good assurance for reliability.

Generally these assumptions were valid for the project; however, where economic considerations are important the choice of where to distribute the PABX would be different. A typical design for a conventional PABX involves a line card which controls 4-8 extensions, a line-shelf that multiplex each telephone onto a TDM bus, and a timeswitch that reorders slots between the in-coming buses and the out-going buses. Distribution of the PABX components can be at the line-shelf, the line-card, or at the extension level in the design. To increase the cost effectiveness of a telephone network-interface future designs would connect telephones to the network though a *Telephone Concentrator*, analogous to the line-card in conventional systems. In an office environment you may imagine the main network being routed along a hallway connecting with a Telephone Concentrator every so often and extensions from these feeding into the local offices.

An opportunity not taken up by ISLAND was to improve the 4kHz bandwidth used for speech transfer in commercial telephone systems. A LAN could easily support an 8KHz, or greater, bandwidth for voice connections and gain a noticeable increase in the quality of the signal.

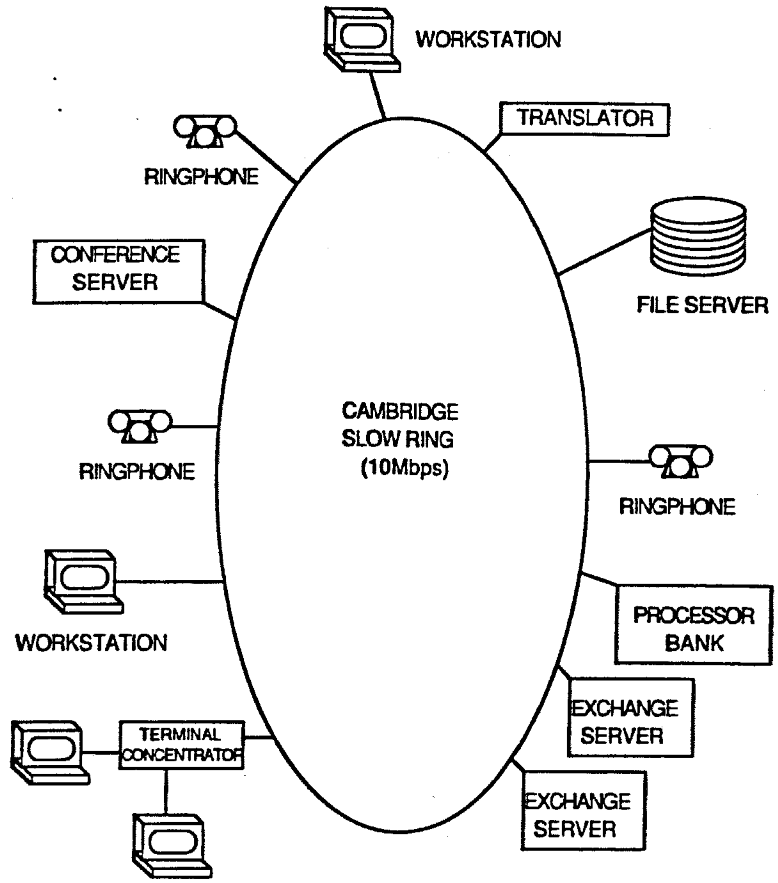

Figure 1: The ISLAND integrated services network

4 Integrated Services

So far a network has been described that is capable of carrying voice and data traffic. This property is little different from the kind of service provided by the ISDN network which is mainly concerned with the end to end communication of voice and data traffic. One of the main advantages of integrating different digital media within a single local-network is that it is possible to combine and interwork with the various media available. An example of such an application is multimedia Hypertext by which a document may be presented to a client in a non-linear form such as a graph. The nodes of the graph may contain: conventional text, graphics, speech, music, stills, or video clips; it is left up to the client as to the order, and how many of the nodes are traversed. Most work in this field is carried out using a workstation incorporating a large graphics display and a variety of expensive support equipment such as computer controlled VCRs, and tape recorders.

Integrated service projects such as ISLAND and Etherphone provide an environment in which media are integrated into the system at the lowest level, allowing the sharing of real-time media storage by workstations in the same way files are shared. In ISLAND a voice dictation machine and voice editor were implemented as a workstation application using the primitives available at the Translator and the ring telephone. In this case a workstation running the application would arrange with the Exchange Service to take over a nearby telephone for its own use.

A similar application can be imagined whereby a linear document contains icons which when opened by a mouse result in the playback of stored digital-speech through a predetermined telephone handset. Or perhaps a document has been electronically mailed for a review and the reviewer wishes to make vocal annotations in the margin of the text. Telephone numbers can also be stored as icons in a document and when activated the system can set up a call to the destination from the nearest telephone.

5 Unison

The Unison project [Clark 86] inspired by project Universe [Leslie 84] and the development of the Cambridge Fast Ring set out to build a land-based network suitable for the provision of integrated services between a number of remote sites. The slower 10Mbs^{-1} Cambridge Rings used in the ISLAND project were already suitable for integrated traffic. Unison realised the potential for using 2Mbs^{-1} ISDN lines as a medium for intersite links. To achieve the interconnection between remote and local rings the faster CFR was used as a high-speed packet-switch in a 19-inch rack thus serving as a hub for communications at each site. A device called a *Ramp* was designed to interface a 2Mbs^{-1} Megastream line and a CFR and another interface called a *Portal* provided connections to lower-bandwidth LANs at that site. A composite network of this kind is a very practical way of providing shared integrated services at a number of sites.

6 The Pandora Project

At Olivetti Research Ltd (ORL) in Cambridge another multimedia project has begun. The project is called Pandora and sets out to provide a video/voice/data communications interface between host computers, initially over a CFR network. The addition of video to such a network allows applications such as videophones and digital video-editors to be an integral part of the distributed computing environment.

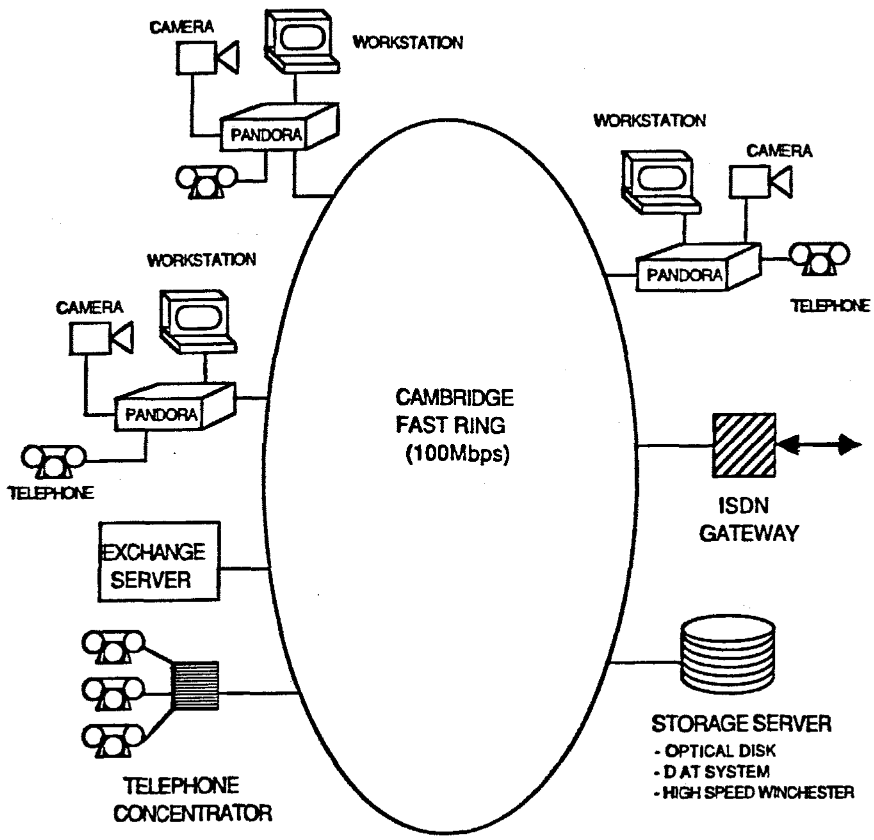

Figure 2: A future Pandora integrated services network

Pandora is a peripheral that can be used with any host computer supplied with a SCSI bus interface. Because Pandora is a peripheral device the handling of real-time media is carried out without loading the host's processor, which in the case of video would severely limit its capacity to execute an application. It can overlay a video frame onto the computers display, for instance inside a window provided by a windowing system, by mixing the computers own composite video signal with the Pandora overlay video. This procedure allows the Pandora video generator to be independent of the way a host graphics adapter is generating its own video. The box also contains a connection to a CFR and a 140MB hard disk, providing communication to other Pandora boxes with the advantage of ATM networking and also an option of local video storage. The audio interface is in the form of a standard BT socket provided on the side of the box, into which any commercial hands-free operation telephone may be plugged.

The first systems will provide black and white pictures with a maximum resolution of 512 x 512 pixels (each pixel being represented by a byte) at a 25Hz frame rate. This would normally imply a transmission rate of 52.4Mbs^{-1}. To make this a more usable figure simple compression hardware has been designed which will reduce this by a factor of 8. Thus each real-time video connection will occupy 6.5Mbs^{-1} of the CFR bandwidth (13Mbs^{-1} in a videophone application). The voice component is telephone quality at 64kbs^{-1} (A-law or μ-law encoded) allowing the possibility of easy connection to the ISDN network. In future Pandora designs colour pictures and the option of high quality audio will also be available.

It is the aim of ORL to produce enough Pandora boxes to assess the value and the usefulness of video applications in a distributed system. It is the view of ORL that video has an important role in the automated office of the future.

7 Conclusion

Multimedia Workstations are likely to be the shape of the personal computer in the 1990s. The ability to handle and manipulate voice and video alongside data is becoming an important feature of the way information is presented in everyday life. Our expectation of the quality and variety in which information is presented is continually increasing. The networks of the future must also allow us to distribute and share this kind of high-bandwidth real-time data.

The ISLAND project demonstrated the advantages of using a single network for integrated voice and data traffic. In the near future the Pandora project will begin to show us the value of video within a local network. Small-packet ATM networks such as the CFR are suitable for the development of multimedia communications and can be used to distribute ISDN services without the need for additional wiring. Furthermore, ISDN lines can be used to provide access to remote multimedia facilities available at remote local area networks. As video compression techniques improve and are implemented as cheap VLSI devices this may allow ISDN and the multimedia LAN world to be integrated into one large network capable of accessing and sharing all types of integrated services.

References

[Ades 87] S. Ades, **An Architecture for the Provision of Integrated Services**, Ph.D. dissertation, Tech. report No. 114, Univ. of Cambridge, September 1987.

[Calnan 87] R. S. Calnan, **ISLAND: A Distributed Multimedia System**, Proc. of Conference Globecom '87, pp744-749, Tokyo, Japan, 15-18th Nov. 1987.

[Clark 86] P. F Clark er al., **Unison — Communications Research for Office Applications**, Electronics and Power, pp665-667, September 1986.

[Greaves 88] D. J. Greaves, A. Hopper, **The Cambridge Backbone Network**, Proc. EFOC/LAN '88, Amsterdam, June 1988.

[Gruber 83] John Gruber, Leo Straczynski, **Judging Speech in Dynamically Managed Voice Systems**, Telesis 1983 Two (Bell Northern Research), pp30-34, 1983.

[Hopper 88] A. Hopper, R. M. Needham, **The Cambridge Fast Ring Networking System (CFR)**, IEEE Transactions on Computers, Vol. 37, No. 10, October 1988.

[Leslie 84] Leslie et al., **The Architecture of the Universe Network**, ACM Sigcomm '84 Symposium, Montreal, Canada, June 1984.

[Swinehart 83] D. C. Swinehart, L. C. Stewart, and S. M. Ornstein, **Adding Voice to an Office Network**, IEEE GlobCom '83 Conference, November 1983.

[Wilkes 80] M. V. Wilkes, D. J. Wheeler, **The Cambridge Digital Communication Ring**, Local Area Communications Network Symposium, Boston, USA, May 1979.

[Want 88] R. Want, **Reliable Management of Voice in a Distributed System**, Ph.D. dissertation, Tech. report No. 141, Univ. of Cambridge, 1988.