

11 An Enabling Infrastructure for a Distributed Multimedia Industry

11.1 Introduction

Advances in technology for telecommunication and new methods for handling media such as voice and video have made possible the creation of a new type of information system. Information systems have become an essential part of the modern world and they need to be made accessible to a very high proportion of the working population. It is therefore important to exploit all the means available for making the transfer of information effective and accurate. In fields such as computer assisted training, multimedia presentation is already well established as a tool for conveying complex ideas. So far, however, the application of multimedia solutions to information retrieval has been limited to single isolated systems, because the bulk of the information required has needed specialized storage techniques and has exceeded the capacity of present day network infrastructure. There do exist special purpose multimedia communication systems, such as those used for video-conferencing, but their cost and complexity separates them from the common mass of computing support.

If, however, *distributed* multimedia systems can be realized, many possibilities for enhanced communication and more effective access to information exist. The key to this new generation of information systems is *integration*, bringing the power of multimedia display to the users in their normal working environment and effectively breaking down many of the barriers implicit in geographical distribution. Now that significant computing power is available on the desktop, integration of voice and video is the next major step forward.

These integrated systems represent a very large market for components and for integrating expertise. It will probably be the largest single growth area for new IT applications over the next ten years. A coordinated set of components, conforming to a common architectural model with agreed interface standards, is required to allow the research and development of prototypes for new applications and to progress smoothly to the delivery of complete multimedia distributed systems. T9000 transputers, DS-Links and C104 routers provide a cost-effective platform on which this infrastructure can be built.

11.2 Network Requirements for Multimedia

11.2.1 Audio Signals

Digital techniques for encoding audio data are well established, and now lie at the heart of the telephone system and the domestic compact disc (CD) player. A number of encoding schemes exist, giving different trades-off between quality, bandwidth and processing costs. Audio support for applications can draw on these techniques and does not pose a major communications problem. However, use of conferencing involving large groups between sites may require a surprisingly high quality of microphone and speaker system to give an acceptable level of reproduction; such environments are often noisy and acoustically complex.

For many purposes, such as remote participation in seminars or discussions, telephone quality speech will be satisfactory. The normal standard for telephony is Pulse Coded Modulation (PCM) [1]. PCM speech will handle frequencies up to 3.4kHz, and is provided as 8k samples of 8bits each per second, or 64kbps. For long distance use, an almost equivalent service can be

provided at 32kbps using the more sophisticated algorithm Adaptively Quantized and Differentially Encoded PCM (ADPCM) [2, 3, 4]. Modern algorithms such as Code Excited Linear Prediction (CELP) can even produce reasonable results at 4.8kbps, but there is no justification for such techniques when communicating with fixed locations on a single site.

Application of ADPCM at 64kbps yields a higher quality speech service, conveying frequencies of up to about 7kHz, which will cover almost all the current requirements. Where higher quality is required (for example, for music or comparative linguistics), one might as well opt directly for a single high quality service, using, for example, CD encoding, in about 0.34 Mbps (stereo). Again, compression will reduce this bandwidth significantly.

The simple PCM encoding is very robust against network loss. The compressed schemes are less so, and the economic balance is probably in favour of compressed data on a moderately reliable network.

11.2.2 Video Signals

The techniques for video transmission are evolving rapidly, with more powerful coding devices giving steadily lower bandwidth requirements. If uncompressed, video information is very bulky, running up to hundreds of Mbps if high quality color is required. Proposed high definition standards, already in use within studios, are even more demanding, with an increase of analogue bandwidth from 15 MHz to 70 MHz and a correspondingly increased digital requirement. However, the information to be sent is highly redundant and great savings can be achieved by compression. Indeed, compressed still images are sufficiently compact to be treated as normal computer data and this section restricts itself to moving images.

There is a significant design choice to be made here: is a moving image to be sent as a sequence of independent still images, or as a progressive representation in which the similarity of successive images is exploited? The latter offers considerably higher compression factors, particularly when motion of objects in the image is detected and exploited. However, this high compression rate is at the expense of greatly increased complexity, particularly if access to the video is to start at arbitrary points.

There are at present three major video compression standards: JPEG [5, 6], MPEG [7, 8] and the CCITT Recommendation H.261 [9, 10].

JPEG (produced by the Joint Photographic Experts Group – an ISO/CCITT committee) provides compression of single images, with compression factors of between 10 and 30, depending on the quality required. There are hardware implementations of JPEG using large scale integration, which give good perceived quality at normal video rates (25 frames per second). A typical PC²⁴-based JPEG card costs about £2,000 at present.

MPEG (produced by the corresponding committee for moving pictures) and H.261 (from CCITT) both exploit interframe coding and can achieve compression ratios of up to 100 or better, depending on the programme material (static material is obviously much more suitable for compression, but quite small scale movements have a large effect on the compression efficiency). However, current implementations are much more complex and expensive, and interfaces with filing systems require research. Current H.261 Codecs cost about £20,000, but much cheaper VLSI implementations are under development.

There also exist other highly effective compression schemes, such as that used in DVI (Digital Video Interactive [11] – a format from Intel) and various fractal-based proposals. However, these suffer from the disadvantage of requiring an expensive compression phase which is slower than real time, ruling them out for many of the intended network applications.

24. PC is a trademark of the IBM Corporation.

All the above encoding techniques have parameters which allow the selection of various qualities of service, the primary parameters being number of points in the image, frame rate and degree of information discarded during encoding. These parameters allow the cost trades-off to be adjusted to meet different quality requirements, so that higher compression might be applied in a general interview, say, than a detailed fine art study. At the bottom end of the range of qualities, there is some competition from rough video provided entirely by software on existing platforms, such as the PC or the Macintosh (Quicktime²⁵), but this low quality material is not a serious competitor for most purposes.

In summary, the most flexible and cost effective technology currently available is that based on motion JPEG. This requires between 2 and 5 Mbps to achieve good quality video from most programme material, although up to 15 Mbps may be needed for guaranteed studio quality. The compressed material is not tolerant of errors, the only effective recovery mechanisms being frame discard and repetition of the previous frame. Future system development based on MPEG or its relatives will offer higher compression ratios at similar costs within five years, but will require low error rate channels.

11.2.3 Performance

Multimedia systems need to be able to capture and present a wide range of media. Some of the media are very bulky, and so present a considerable challenge to network and operating system designers. The most demanding requirements come from isochronous media, such as audio and video, since they have fixed timing deadlines.

The network requirements can be characterized in terms of the necessary bandwidth and end-to-end delay, and by the acceptable variation, or jitter, in the delay. Media transmitted in their raw form also show different tolerance to loss of data samples, but as increasing use is made of powerful compression techniques, data loss becomes correspondingly less acceptable.

To a considerable extent, the demands can be matched to the available network resources by adjusting the quality of reproduction offered. Both audio and video remain usable for many applications through a wide range of qualities. Low bandwidth allows understanding, but higher bandwidth increases quality, pleasure and impact of the presentation. The particular demands from a range of media are summarized in Table 11.1.

25. Macintosh and Quicktime are trademarks of Apple Corp.

Table 11.1 Bandwidth Requirements

Medium	sample size (bytes)	repetition rate/duration	uncompressed rate	compressed rate
Text Page (A4)	5k	1 sec	40k bps	10k bps
Colour image (640 x 512 pixel)	320k	3 sec	800k bps	90kbps
Audio (3.4kHz)	1	8k/sec	64k bps	4.8k bps
Audio (7 kHz)	1.5	16k/sec	200k bps	64k bps
Video (high resolution)	320k	25/sec	65M bps	7.5M bps
Video (low resolution)	20k	15/sec	2.5M bps	200k bps
Videophone (static subject)	20k	15/sec	2.5M bps	64k bps

11.3 Integration and Scaling

In addition to the requirements of the media themselves, the need to integrate them into a single system must be considered. In multimedia applications, presentation of a number of different media needs to be coordinated. Typically, this implies a need for some form of general distributed platform, providing efficient and flexible communication for control and synchronization mechanisms. Modern object-based platforms can meet these requirements in a flexible manner.

Until recently, most multimedia computer systems were constructed on the basis of computer control of essentially analogue systems. For example, interactive video systems generally consisted of analogue videodisk equipment controlled by, and sharing display facilities with, a personal computer. Such analogue systems do not scale well. Analogue video networks are difficult to maintain, manage, and share between different applications.

Developments in the technology available have now made possible the construction of equivalent digital networks, and digitally-based multimedia systems can now be constructed. This opens up the possibility of multi-service networks (both local and wide area) which can convey a range of multimedia information types on a single network, giving economies of scale, flexibility and ease of management. In this environment, site-wide distribution of audio and video information, integrated with traditional computer data and control, becomes a realistic proposition.

11.4 Directions in networking technology

From the computer user's perspective, network developments over the past ten years have been dominated by the increasing coverage and performance of the local area network; there now need be few barriers to the sharing of text, data and program within a site. For the more demanding media, particularly for video, current networks can support single user demonstrations, but not the activities of a realistically sized user community.

For example, the bandwidth requirements listed in Table 11.1 indicate that a number of existing technologies would be able to support the requirements of a single multimedia station (e.g. serv-

icing a small office or conference room). Such requirements probably do not exceed 20 Mbps each way in total and current ring and bus technologies (such as Fibre Distributed Digital Interface (FDDI) [12], FDDI-II and Distributed Queue Dual Bus (DQDB)) all have the necessary raw capacity – although the capabilities of their routers and the ability to reserve bandwidth are more questionable.

In reality, however, a building (or group of buildings) will require the parallel operation of many such stations. The University of Kent, for example, has some 125 teaching rooms registered for AVA provision. Even if only 20% of these were using multimedia at one time, the total bandwidth requirement would be almost a Gigabit per second, beyond the capabilities of any of the current ring or bus networks. Something with an order of magnitude greater capacity is needed.

Fortunately, a quiet revolution has been taking place in the wide area networks, exploiting the power of fibre optic transmission and providing the basis for a telecommunication infrastructure of enormous capacity. The aim is to produce a truly integrated network that is able to carry all types of traffic, from the isochronous data of digital voice and video to the bursty packet traffic produced by computer applications, using a single underlying network technology.

The approach being taken is based technically on the CCITT Recommendations for the so-called Asynchronous Transfer Mode and the resulting networks are called ATM networks. They use the efficient switching of small fixed-size (53 byte) packets to provide a combination of high speed to any individual user and simultaneous service to large numbers of customers. The small fixed-size packets are called cells and their use facilitates the multiplexing of traffic that is sensitive to delay jitter with traffic that is not. ATM systems designed for telephony are expected to operate at speeds of 155Mbps and above. This implies a packet switching rate of over a quarter of a million packets per second from each link into a single switch. The distribution of the switching function means that the total switching and transmission capacity is not limited to that of any individual link or switch; the system scales up naturally like the telephone system and does not suffer from size limits like a LAN.

11.5 Convergence of Applications, Communications and Parallel Processing

11.5.1 Multimedia and ATM

The capabilities of the new ATM networks are well matched to the requirements of distributed multimedia systems. ATM networks operate at a high enough speed to support all the types of information wanted, and give the flexibility needed to share information between many users.

Recent advances in compression technology also affect the situation by reducing the peak requirements of audio and video to the point where they can be handled and stored in conventional desktop systems. The ATM networks cope well with the varying load presented by compressed data streams.

The future therefore seems clear:

- there will be an increasing penetration of ATM technology as the networking solution of choice, both in the wide area and as the local infrastructure;
- there will be a major expansion in the use of multimedia technologies for integrated communication and information retrieval both within and between organizations.

As an example, the UK academic sector has recently initiated the SuperJanet project to provide a new generation of wide area infrastructure capable of supporting multimedia applications, and a number of multimedia research groups are planning to make use of it to extend their local facilities.

11.5.2 ATM on your Desk

However, the currently available ATM equipment is primarily aimed at the telecommunication carriers. The real benefits of an integrated ATM network become apparent when the flexibility of the mixture of ATM based services is delivered to the end user's desk [14]. This implies a need for local area ATM solutions, and for ATM compatible end user equipment.

In the multimedia architecture described in this chapter, the further step of using ATM cells for communication within the workstation is proposed. Multimedia activities produce a great deal of movement of data and the same considerations regarding the transmission and switching of this data pertain within the workstation as in the local and wide area. It therefore seems sensible that a similar solution should be applied.

At present, however, ATM switches on offer commercially are disappointingly expensive. Prices of £75,000 for a small switch are typical and this renders – for the time being – their use uneconomic for supporting the above scenario.

11.5.3 Transputers and ATM

Starting from the need for the flexible interconnection of parallel processing elements that is required to produce parallel computers, a style of architecture has emerged in which separate processing elements are interconnected by communication links. Some of these designs are packet based, and the best known is the INMOS transputer. In the T9000 range of transputers, INMOS have chosen a style of communication which is similar both in general philosophy and in technical capability to the ATM networks. The rationale behind this decision is broadly similar to that which led the CCITT to choose ATM for the Broadband Integrated Services Digital Network (B-ISDN).

The INMOS choice has the happy consequence that it will be possible to construct systems which, with a minor amount of technical 'glue' to make the necessary detailed adaptations, carries the same high level view of ATM based communication from the wide area into the local processing component within multimedia devices. Integrated multimedia networking becomes possible at reasonable cost (with the cost of an ATM switch being reduced by at least one order of magnitude, probably more).

11.5.4 Convergence ... and an Opportunity

This three-way convergence of application requirements, telecommunications standards and parallel processing technology represents a real opportunity for progress; all the important pieces are becoming available now (1993). It is therefore the right time to seek to define standards so that the necessary components can come together from different vendors to form a single family of compatible products.

11.6 A Multimedia Industry – the Need for Standard Interfaces

The development of multimedia applications depends on the availability of suitable products. Broadly based applications can only be constructed easily if the necessary components are offered by a number of suppliers in a form which is simple to integrate and to configure to meet the specific needs of the user. The detail of the configurations needed on a desktop and in a meeting room will be different; so will the configurations needed by the author of training material and users of that material.

The solution to this problem lies in a modular approach, based on the definition of a small set of key interfaces between components. The interfaces are crucial because the exact packaging

of functions and the power of the components themselves will evolve as the technologies develop. The interfaces between components, however, are relatively stable and allow the construction of systems using components from different sources and with different levels of technical sophistication.

Some interfaces need to be common to all components; others are more specialized. Universal interfaces are needed for the control and management of the components – particularly for:

- the control of communication and information storage and retrieval, so that common tools and common user paradigms can be applied across the full range of media;
- those enquiry functions which allow each component to determine the capabilities of any others with which it interacts and so take account of the changes and limitations of the configuration in which it is placed.

At a more specific level, agreed interfaces are needed for each of the media types so that, for example, all audio or all video components can interwork successfully.

To a large extent, these interfaces can be based on internationally accepted standards, but in some of the areas being addressed such standards do not yet exist, or the way they are to be combined is not fully defined. There is an urgent need, therefore, for the involvement of a broad range of potential component suppliers and system or application developers. What is required is the minimum of technical work for the definition of a profile for multimedia use of ATM standards. It would provide both an architectural framework which identifies the interfaces needed and a portfolio of references to related interface specifications covering the full range of multimedia requirements.

The agreement by the Industry to such a profile would provide a firm basis for the development of new applications using distributed multimedia systems and would be a major input of practical experience into future formal standardization.

11.7 Outline of a Multimedia Architecture

11.7.1 Components, Stations and Sites

The basic building block of the architecture is the ‘component’. Components will either handle a piece of equipment (such as a microphone, video camera, display, disk, ...) and/or carry out a function such as encryption or compression. Each component will contain one or more processors. The most important feature of the architecture is that the components will communicate with each other in a single universal fashion by the transmission and reception of ATM cells. ATM cells will be used to carry the media, to carry control information and to carry signalling information (i.e. the commands for organizing the pattern of interconnection between the components). Components are viewed under this architecture as atomic devices – i.e. communication mechanisms *within* a multi-processor component are specific for that component and would not follow ATM standards.




Figure 11.1 Hierarchy of interconnections

Three regimes of ATM interconnection between components are envisaged (see Figure 11.1). The first regime is that of a 'station'. A station is a set of one or more components that can be considered to act together – i.e. they are either all switched–on or all switched–off, and the transmission of cells between them can be considered to be error–free. The second regime is the local interconnection of stations – a 'site'. This is a regime in which transmission delays will be short, and error–rates will normally be very low. It must cope with stations varying their status between inactive and active. Such a regime may use synchronous or asynchronous transmission. The third regime is that of the 'wide–area', in which it is assumed that the full application of CCITT standards will be the norm – i.e. quality of service, policing, tariffing, etc.. In this regime, synchronous transmission will be used. Within all three regimes there will be components that are used for switching cells.

The architecture has three main areas that require agreement on standardization:

- specifications of how different types of media are to be carried in cells. As far as possible this will follow international standards – i.e. use of ATM adaptation layer, use of standard encoding for voice (CCITT G.series Recommendations) and video (JPEG, MPEG, H.261), etc.;
- specifications of how components are to be controlled. This will have two parts: a general scheme of control and realization of this scheme for particular components;
- specifications of the manner in which signalling is to be carried out – i.e. how connections are to be created and removed within the three ATM regimes. Clearly in the wide–area regime this will be determined by outside authority.

11.7.2 Example Components

The hardware requirements to support audio and video revolve around a family of interfacing components delivering and controlling the media. In each case, the information flows from the network as a stream of ATM cells. The interface decodes and decompresses the information, converts it to analogue form and passes it to the display device or audio system – see Figure 11.2. For input devices, the process is reversed.




Figure 11.2 A Component

The display (or capture) can be performed by connecting existing audio– visual devices, although more integrated solutions will appear as time goes by. The controller element can be a specialization of a single design for the whole family, but the decoder (or encoder) is specific to the medium being supported. The network interface can be expected to have two variants: one for direct connection as a station to the site network for use by isolated devices and a largely vestigial one for use within a station for connecting the different components supporting multiple media.

The following components are considered to be basic to the architecture:

- video capture (still and full motion) including compression;
- display including decompression and input device handling (keyboard, mouse, etc.);
- audio input and output (including compression);
- bulk media storage;
- encryption;
- switching both within a station and in the local area;
- access to other networks and other communications technologies.

This is an initial list of areas to be covered and it will grow as the industry develops. The central component is a very small ATM switch (on a single card) to integrate a local cluster into a station – a single C104 router and a minimal controlling processor may be all that is required.

11.7.3 Example Station

Using the components outlined above, a suitable multimedia station for an office or conference room can be constructed from modular components. The size of the modules will be determined largely on economic grounds, relating to processing costs. For example, it may be attractive to provide both audio input and output in a single card, but to separate video input from output.

A typical configuration would always provide a hardened network interface for the station as a whole, an enclosure and power supply, together with a small integrating switch. The switch

would be used by components specific for that station (e.g. a full-duplex audio card, a video input card, a video output card and a control terminal interface, also providing OHP tablet output) to share a single outgoing ATM link to the site multiservices network – see Figure 11.3.




Figure 11.3 A Station

11.7.4 Example Site

The campus of the University of Kent at Canterbury provides a typical ‘site’ for such a multiservice network. Geographically, it is a compact single area with most of the major teaching buildings falling within a circle of 500m radius. There are some outlying locations, but none are more than 2 Km from the centre. The campus is crossed by public roads, but the University has ducting under them.

Fibre optic links have been installed throughout the campus and services are being migrated onto them. The new links provide for an FDDI-based backbone and a number of distribution links to the Ethernet segments in individual buildings. Provision varies from 12 fibres per link in the central ring to 8 or 4 fibres on the distribution spurs. The central part of the campus, together with the fibre network, is shown in Figure 11.4.




Figure 11.4 The Fibre Routes at UKC

These optic fibres (or, at least, the physical channels in which they are laid) would form the backbone for a 'site' ATM network – once the low-cost distributed multimedia industry described in this chapter and enabled by T9000/DS-link/router technology comes into place. The logical structure of a possible (initial) UKC site is shown in Figure 11.5.




Figure 11.5 A Site

11.8 Levels of conformance

Combination of modular components can be viewed at a number of different levels. The more detailed the specification used, the lower the integration cost, but the more limited the field of application. It can therefore be worthwhile to identify different levels of conformance to the interface specifications.

One can distinguish:

- an abstract statement of the media types, the processing components and the interfaces and data flows between them. It is the essential minimum set of agreements necessary for system integration to be possible, since it includes the agreements on data types and interpretations needed to have a common understanding of how to process and represent the various media. However, it does not commit an implementor to any particular communication technology or physical packaging. Using the terminology of the international standards for Open Distributed Processing (ODP) [13], this corresponds to ODP 'information and computational specifications';
- a statement of how the various interfaces are to be realized, giving the detailed constraints on implementation in a particular environment. This corresponds to ODP

‘engineering and technology specifications’. Several different solutions may be needed to support different kinds of integration. Of particular importance are solutions to:

- *network interconnection*: specifying the formats and protocols that are to apply between two systems on a wire or a fibre. This form of specification does not constrain the internal structure of the systems and is the minimum requirement for the construction of distributed applications;
- *physical packaging*: specifying the form and interconnection requirements for a system component at, for example, the card level. Widely accepted standards for particular computer families, such as the format for PC cards and buses fall into this category;
- *software interfaces*: specifying the interface to device drivers and presentation management systems at a language level. These specifications should be obtained directly by selection from established industry practice, rather than creation of new specifications. Support for specific software environments such as MS-DOS/Windows²⁶ or UNIX/X-Windows²⁷ falls into this category.

This framework then allows integration to take place at many different levels, but within this structure all players are expected to conform to the abstract specifications. All suppliers of communication components are expected to conform to one of the agreed communication specifications. All suppliers of, for example, PC cards are expected to conform to the specifications for the physical, bus and device driver specifications for the machine range. A supplier of a video display card with an integral network interface might need to conform both to the networking and the subsystem interfaces (see Figure 11.6).

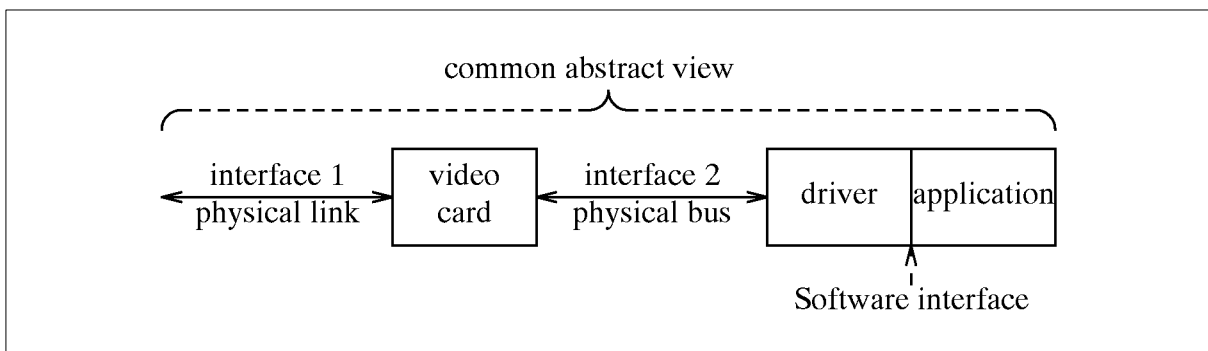


Figure 11.6 Possible conformance points

11.9 Building stations from components

As well as the abstract standard for the architecture, hardware standards such as methods of transmission of cells, particularly within the station, board standards, etc. have to be specified for the concrete realization of components.

For the transmission of cells within a station there are broadly two possibilities: use of a standard bus or use of point-to-point links in conjunction with routing between components. Use of a bus has two major disadvantages. First, it would put a non-scalable resource at the centre of the station, which would, moreover, be a shared resource whose properties would have to be taken into account when various combinations of components are integrated together in a station. Second, there are a large number of possible bus architectures that might be chosen.

Links do not suffer from these disadvantages; they are scalable and they exploit the same interconnection model between components as has already proved effective at higher levels (between stations and between sites). The approach taken is thus logically coherent.

26. MS-DOS and Windows are trademarks of the MicroSoft Corporation.

27. UNIX is a trademark of AT&T Bell Laboratories and X-Windows is a trademark of MIT.

A final question that must be addressed in the definition of the architecture is its relationship to various existing workstation architectures. The two main architectures to be considered are the Unix-based stations and the PC-based stations. The Macintosh architecture has strong claims for consideration, but certainly runs third to the others. Interfacing between workstations and the multimedia architecture is principally in the areas of the display screen, control of the multimedia components and access to files. At a minimum, interfacing through X-windows and/or MS-Windows, via a simple RPC mechanism and via ethernet, will be required.

11.10 Mapping the Architecture onto Transputer Technology

T9000 transputers, DS-links and C104 routers are well-suited for the construction of low-cost ATM networks – detailed technical analysis to support this claim is presented in chapter 10. On top of this, INMOS have defined a board technology for the construction of modules. This technology defines a board format called the H-TRAM for small boards that plug into motherboards. Thus, if most of the multi-media components are built as H-TRAMs, they could be used with different motherboards to fit a variety of situations. Motherboards dealing with switching and interfacing functions are likely to be built for all the popular bus standards.

For communication and switching between components within a station, T9000 technology provides the necessary means of integration directly – without further development. ATM cells can be conveyed directly over DS-links, routed through a small C104 network (one chip will generally be sufficient).

The use of transputer parts between stations in the local ATM regime and the interfacing to wide area ATM will require the development of specialized chips. For the local area regime, a part is required that will allow INMOS links to be carried over distances of up to a few hundred metres. This part must also provide guaranteed immunity of the component at one end from any type of failure of the component at the other end. This isolation is necessary because the different stations in the local area belong to different people and may be powered up or down (or reinitialized) independently of each other and of the switching and communication components.

To interface to the wide area, a part suitable for interfacing a T9000 processor to a synchronous link running at up to 155 Mbps is required. However, this is a peak speed and represents the loading of a multiplex of many user activities. It is therefore possible to distribute it immediately onto a number of DS-links in all but the most pathological congestion situations, where higher level recovery can be expected to take place. The initial 155 Mbps serial link interfacing requires moderately fast hardware, but is well within the capabilities of available components.

Some preliminary investigation of these requirements has been made and it is felt that both the local area and the wide area problems can be solved by an adaptor constructed using electronically reconfigurable programmed logic arrays rather than custom designed chips. However, the T9000 link engine is expected to be available as a semi-custom library component, allowing the creation of multisourced low-cost components as the market grows.

REFERENCES

- [1] CCITT Recommendation G.711: Pulse Code Modulation (PCM) of Voice Frequencies.
- [2] CCITT Recommendation G.721: 32kbit/s Adaptive Differential Pulse Code Modulation (ADPCM).
- [3] CCITT Recommendation G.722: 7kHz Audio Coding within 64kbit/s.
- [4] CCITT Recommendation G.725: System Aspects of the Use of the 7kHz Audio Codec within 64kbit/s.

- [5] ISO/IEC 10918: Information Technology – Digital Compression and Coding of Continuous Tone Still Images (JPEG).
- [6] Wallace, G.K., "The JPEG Still Picture Compression Standard", CACM, vol 34, pp. 30–44, April 1991.
- [7] ISO/IEC 11172: Information Technology – Coding of Moving Pictures and Associated Audio for Digital Storage Media (MPEG).
- [8] Le Gall, D., "MPEG: A Video Compression Standard for Multimedia Applications", CACM, vol 34, pp. 46–58, April 1991.
- [9] CCITT Recommendation H.261: Video Codec for Audiovisual Services at p*64k bit/s.
- [10] Liou, M., "Overview of the p*64 Kbit/s Video Coding Standard", CACM, vol 34, pp. 59–63, April 1991.
- [11] Ripley, G.D., "DVI – a Digital Multimedia Technology", CACM, vol 32, pp. 811–822, July 1989.
- [12] Ross, F.E., "An Overview of FDDI: The Fibre Distributed Data Interface", IEEE J. on Selec. Areas in Commun., vol 7, pp. 1043–51, Sept 1989.
- [13] ISO/IEC 10746: Basic Reference Model of Open Distributed Processing, Part 2: Descriptive Model and Part 3: Prescriptive Model.
- [14] Hayter, M.D., and McAuley "The Desk Area Network" Operating Systems Review vol 25, no. 4, October 1991

