THE TECHNOLOGY OF BROADBAND NETWORKS

Robert A. Mercer

1. Introduction

In this paper, we explore the candidate technologies and associated network architectures being considered for the Integrated Broadband Network (IBN).[1] We begin in Section 2 with a definition of an IBN based on the applications that have been identified and which are discussed more fully elsewhere in this volume. Section 3 discusses candidate network architectures, comparing them with the architectures currently in use. This leads to the identification of architectural changes that will probably be necessary in implementing an IBN, and to other issues requiring further consideration.

In Section 4, we identify and discuss technology options in four principle areas:

- customer premises equipment (CPE) and distribution media including Local Area Networks (LANs), terminals, building cabling and wiring, etc.;
- network access (often called "local distribution");
- switching; and
- interoffice transmission.

They are depicted in Figure 1. The CPE shown in the premises portion of this drawing should be understood to include Local Area Networks, which consist of devices (computers, printers, etc.) interconnected by a transmission medium such as coaxial cable, optical fiber, or twisted copper wires.

Integrated Broadband Networks: The Public Policy Issues / Martin C.J. Elton (Editor)

Figure 1.
Components of the IBN

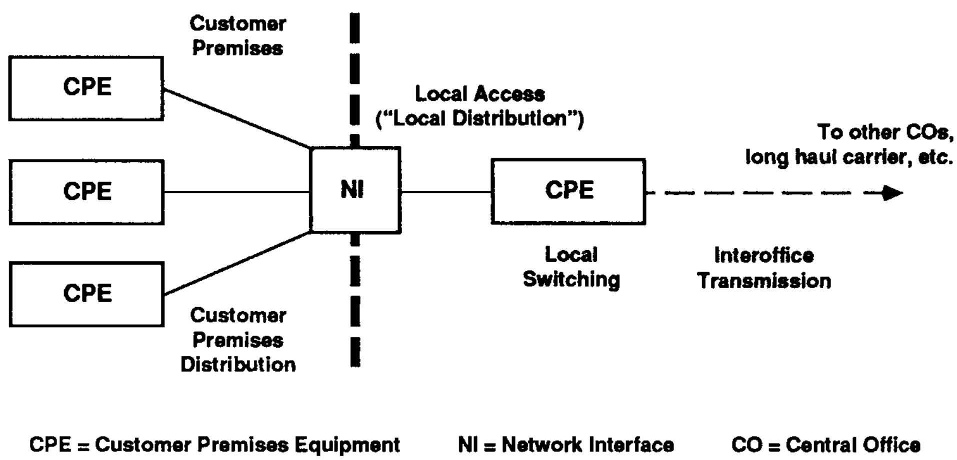

Given that the IBN is quite different in form and capabilities from today's networks, it is important — as a precursor to the discussion of associated policy, social, and regulatory issues — to understand how IBN might evolve and over what period of time. Section 5 explores this matter, both in terms of technical and operational issues, the latter leading in turn to a discussion what kinds of organizations might operate an IBN.

Section 6 presents our conclusions as to the nature of IBN and the evolutionary path needed for attaining it.

2. Definition of IBN

Rather than address the issues concerning the social or economic utility of IBNs, it will be taken as a given that there is a need to develop networks which can transport information in digital form at speeds much higher than today's networks. This means local distribution rates of hundreds of megabits per second (Mbps) and interoffice rates in the gigabit (10^9), or conceivably terabit (10^{12}) per second range. Applications creating this need might include:

- television program distribution, both in conventional form and in new signal formats such as those being conceived for high-definition television (HDTV) and extended quality television (EQTV).

- other video applications such as video conferencing and video shopping;

- LAN interconnection and, possibly, the eventual replacement of LANs themselves; and

- general networking of ultra-high-speed computers for information intensive applications such as CAD/CAM.

It is assumed that at least some of the applications listed above require *switched* transport.

We define an IBN as follows: a public switched or dedicated network capable of supporting the integrated transport of voice, data, and video information in digital form at user-accessible rates in excess of 100 megabits per second (Mbps). To better grasp the significance of the definition, let us contrast it with the networks utilized today. They involve a mixture of analog and digital transmission. They are not integrated to a great extent; rather, different applications require different network services. While users can in some cases obtain transmission rates on dedicated circuits as high as 45 Mbps, the rate more typically available is 1.5 Mbps. As far as digital switched services are concerned, rates are limited to 56 kilobits per second (kbps). In fact, most users employ modems on analog dial-up circuits with a practical data rate limitation of less than 10 kbps! IBN thus represents a major step forward, increasing by a factor of ten thousand or more the rate of switched digital services.

Let us consider the implications of the various terms in the IBN definition. First, by "public" we mean only that the network extends beyond a single customer premises and involves the use of transmission facilities available from a public provider. We do not mean that the network is of necessity a shared-use network, although that will normally be the case.

We restrict our attention to public networks because we believe the key issues being dealt with here are of primary interest in that realm. We emphasize that most of the technology discussed here can be applied equally well by a single user to create a premises-contained network, as opposed to acquiring IBN-based services from a network provider. Furthermore, we do not want to exclude the premises portion of a public network from consideration since there are several technology questions with significant policy and economic implications. We will discuss these issues later.

Next, consider the modifiers "switched" or "dedicated". From a technological point of view, switched networks are of greatest interest. The IBN problem is much less interesting, and in many respects already solved, if it has to deal only with dedicated high-speed transmission links between premises. On the other hand, we recognize that there are policy issues involved even in dedicated networks, such as, who can provide transmission services and under what terms and conditions? We suggest that the IBN concept should embrace both types of networks, even though most people's notion of IBN involves switching.

IBN involves the notion of integration. This is commonly taken to mean that all information can be efficiently transported in a common format over a single network regardless of the application supported (voice, data, video, etc.). It is not a given that a single network can accomplish this cost effectively. This issue is still very much alive, for instance, in the deployment of Narrowband Integrated Services Digital Networks.[2] The subsequent discussion deals with this issue in two key areas; interface standards and switching. At worst, one can view the *I* in IBN as indicating that all applications are being cared for in one plan, even though the plan may entail more than one network fabric.

We use the term "transport" throughout this paper. Consistent with the practice emerging in the industry, the term is synonymous with the movement of information on a switched or dedicated basis. As a subset of transport, we utilize the term "transmission" to indicate movement of information over a single facility from one point to another.

We have restricted the definition of IBN to the transport of digital information, in accordance with the predominant thinking in the industry. It should be noted that for at least one of the principal applications of IBN, entertainment video, there is a body of thought that analog transmission is the preferred choice, because of the massive number of television sets that use an analog format. For such sets, a codec would be required for conversion from the digital format, and it is anticipated in some quarters that there will be substantial user resistance to incurring the cost of adding such a device.[3].

IBN is initially assumed to be restricted to digital transmission. If digital television program distribution does not gain marketplace acceptance for reasons of convenience or economics, then it will be outside the scope of IBN as initially defined. We suspect that this omission would deal a potentially fatal blow to the prospects for IBN, for at least in the early going, television appears to be the principal motivator and likely determinant of success. In that case we would expect the definition of IBN to broaden to encompass analog transmission as well. A hybrid arrangement is already being discussed in which more than one fiber would be connected to each customer premises, with one or more fibers carrying analog signals for video services, and one carrying digital information. Another possible evolution involves sharing the wideband capacity that a single fiber provides between digital and analog signals.

Finally, we have assumed that IBN involves transport speeds of 150 Mbps or above because of the requirements imposed by various applications. This would be in line with the current industry concepts of IBN. This does not mean that there will be no coupling of access links operating at less than this rate, but only that access to at least that rate is required. As noted previously, the upper limit on transport rates in the interoffice portion of the network, and possibly even in the access links for some advanced applications, might be several gigabits per second (Gbps).

Several other attributes of an IBN are commonly assumed.[4] First, an IBN should be capable of interfacing with existing LANs and other terminal equipment. Second, because existing applications which are to be integrated via IBN involve a vast range of transport rates, an IBN should allow the dynamic allocation of available bandwidth under user control, rather than having a fixed-rate channel definition scheme. Further, the switching technology utilized must include a multi-rate capability. Third, delay caused by queuing and processing in the network should be low. Fourth, due to limitations on the amount of processing that appears to be possible if high throughput and low delay is required, IBN should provide relatively little error detection and correction. That function should be left to higher-level devices. Finally, IBN should support current data communications protocols through simple and efficient interfaces which maximize throughput.

3. IBN Architectures

Because entertainment video is considered to be a principle application of IBN, the architecture and technology discussions in the next two sections deal with both telecommunications and entertainment video networks as they exist today. The current telecommunications network infrastructure consists of four discrete components: customer premises equipment and distribution media, network access links connecting the premises to the serving central office (CO), local switching, and interoffice transmission and switching. These components apply to entertainment video as well, although there is no true switching function since signal distribution is done in a broadcast mode.

Customer Premises Equipment

The typical customer premises arrangement takes one of two forms, or a mixture of both. The first form consists of twisted wire pairs fanning out from a network interface (NI). Connected to each wire pair individually or on a multiple-tap basis are one or more devices: telephones or data communications equipment such as terminals and PCs. The NI is the point at which the access portion of the network terminates near the entry point of the customer premises. It is typically little more than a device that allows the mechanical interconnection of inside and outside transmission media. It also serves to protect the premises against voltage surges due to lightning strikes or close proximity to high-voltage power lines. This is the predominant arrangement in residences and small businesses. The same arrangement is used for CATV except that the use of multiple television receivers or other video equipment normally requires a device that splits the signal onto multiple paths.

In the second arrangement, typical of most medium to large businesses, individual terminals are linked together through a LAN, or several interconnected LANs.[5] A LAN consists of a shared-use medium, such as a coaxial cable or fiber optics cable, or a twisted wire pair, capable of supporting high-speed digital communications between terminals, work stations, and other devices. In order to achieve high data rates with reliability, and to simplify the protocols utilized, LANs have a limited geographic range — typically not more than a few hundred meters, but extending as far as 100 kilometers for some types of LANs. LAN transmission rates are typically in the range of 0 to 10 Mbps but are rapidly pushing towards 100 Mbps or higher as fiber optic LANs develop.

LANs commonly use a bus or ring topology. Some PBXs, which were originally limited to voice transmission, now provide LAN capabilities. They utilize a star arrangement in which each terminal is connected to the PBX over a separate link.

Several LAN topologies and corresponding protocols have been standardized and are supported by a number of manufacturers. Whatever form of LAN is used, information whose destination is outside the LAN passes through a gateway to either another LAN or to the NI, depending on the destination of the information. The gateway may be complex or simple, depending on whether it is connecting two networks with the same or different protocols. At one extreme, a simple bridge between two identical LANs does little more than provide a transparent "pipe" between them. At the other extreme, a gateway that connect a LAN to a "connection-oriented" network can require an additional protocol level to function properly.

Network Access

The network access portion of the telecommunications network has a star topology, fanning out from the CO to the premises of the customers served by that CO.[6] Historically, the access links to each premises have consisted of individual pairs of twisted copper wires; the wires are twisted together to minimize electromagnetic interference from other signals. The original pattern has evolved into a more common arrangement in which the access portion of the network is divided into two components. In the so-called "feeder plant," high-capacity links fan out from the CO and terminate at serving area interface points (SAIs). In turn, the "distribution plant" portion forms a second star around each SAI, and connects the individual premises to the SAI. A carrier system is often used in the feeder plant. It multiplexes a number of access links together onto a single high-speed link for transmission to and from the CO. The distribution plant still consists of wire pairs to residences and small businesses, while carrier systems may connect directly to the premises of large businesses which require large numbers of circuits and/or high-speed digital access. Although there may eventually be some switching functions performed at the SAI, the cross-connections at the SAI are likely to persist for the foreseeable future. This being the case, the access network is logically a single star, even though physically it appears as a double star, because each premises has, in effect, a dedicated circuit all the way to the CO.[7]

A recent, but rapidly accelerating trend is the use of fiber optics in the access portion of the network. Fiber is deployed most frequently in the feeder plant but is now being extended to customer premises in several trial service offerings. In a typical configuration used to distribute television signals, one digitally-encoded television channel is carried on each of a large number of fibers which traverse the feeder plant portion of the access network. One to three fibers run from the SAI to each customer premises. The customer is able to select up to three channels, for simultaneously viewing. This is done by directing a simple cross-connect switch at the serving area interface to select the multiple feeder fibers it should connect to each of the distribution fibers. In early trials, the necessary signaling is being done over a separate wire pair from the premises to the SAI, but as the technology evolves, it may be done over the fiber itself.

A second use of fiber optics in the access network is to support high-speed digital transmission of voice and data for large users. These users need the high capacity 45

Mbps rate (DS-3) that is derived from a fiber optics link. Although today's applications involve high-speed transmission of multiple digitized voice and data signals and the distribution of entertainment video, projections are that such a system will become cost effective for new installations of voice service within the next five to ten years; in some cases perhaps sooner, in light of the fact that the access network plans of most telephone companies now require two distribution links per new home.[8]

The "access" portion of a cable television network utilizes a tree arrangement in which trunks fan out from a headend (equivalent to the telecommunications CO) and are split into a number of branches. Individual subscribers are tapped into the branches. In many respects, the physical topology resembles the star arrangement used by telephone companies. The difference is that for entertainment video, the same signal is carried throughout the distribution area, with program selection occurring only at the receiver. This contrasts with the telecommunications model in which each subscriber link is carrying different information.

Local Switching

Let us consider the local switching portion of the telecommunications network. Here, information generated at the premises arrives on the access links (outbound traffic) and is directed either onto another access link, in the case of an intra-office call, or onto interoffice circuits. Likewise, inbound traffic is connected to the appropriate access link for delivery to the destination premises. For switched services, the appropriate interoffice link (for outbound traffic), or access link (for inbound traffic), is determined by the addressing information provided by the originator of the information. For dedicated services, the local office does not switch the call at all, but it still may provide various static or low-speed cross-connection mechanisms, transmission equipment such as multiplexers and signal shapers, and access points for various test capabilities.

Current switching systems are for the most part limited to circuit-switched analog circuits, augmented in some cases by digital switching up to a rate of 56 kbps. Circuit switching means that a circuit is established from the originating customer to the dialed customer and held for the duration of the call.

Such fixed circuit utilization is not efficient for most data applications where the transmission of information is "bursty" in nature. As a result, packet switching is being increasingly deployed for public use by local exchange companies, interexchange carriers, and value added network (VAN) providers (Telenet or Tymnet), supplementing the extensive use in private computer networks over the past two decades. In packet switching, which is largely based on the well-known international standard X.25, information associated with a given data session is transmitted in segments called packets. Packets from different sessions are interleaved over a given transmission link to maximize its utilization. Access links to such networks normally operate at rates of 56 kbps or below. A new technology appearing on the scene, Frame Relay, will utilize much faster switches, and allow access at rates as high as 1.5 Mbps.

Interoffice Switching and Transmission

Finally, the CO provides interconnection between the network access and interoffice portions of dedicated circuits. As mentioned previously, such circuits are available at digital rates of up to 45 Mbps (DS-3), but 1.5 Mbps is much more common.

In the case of entertainment video, only dedicated cross-connects are made at the headend, as there is no switching per se. On an infrequent basis, different program signals may be cross-connected to a given channel for distribution to customers. These infrequent changes in cross-connections are sometimes referred to as "channel switching".

Consider finally the interoffice portion of the telecommunications network. Historically, it has utilized a backbone network consisting of a hierarchical arrangement of interconnected switches, augmented by direct connections between switches at all levels of the hierarchy as justified by traffic engineering considerations. These direct connections move the network towards a "mesh" (fully-connected) topology. With the divestiture of the Bell Operating Companies (BOCs) from AT&T in 1984, and the resulting distinction between networks which provide exchange or metropolitan-area services and those which provide long distance (inter-area) service, two different hierarchies have become evident. In metropolitan areas, the BOCs and other exchange carriers (hereafter collectively referred to as Local Exchange Carriers, or LECs) typically employ a two-level hierarchy consisting of COs at the first level, and a second level of tandem switching which interconnects the COs. In some cases, there is a third or "principal tandem" level. Again, notwithstanding this hierarchy, the extensive use of direct CO-to-CO connections, where warranted avoids a stringent hierarchical structure and moves the network towards a mesh topology.

Around the time divestiture was occurring, AT&T introduced a new concept called dynamic non-hierarchical routing (DNHR).[9] In effect, many of the switches in the upper levels of its former hierarchical network are now connected in a mesh topology, and routing is determined in realtime according to existing traffic conditions. In this portion of the network there is no longer a hierarchical structure. Thus the AT&T network today employs a mixed hierarchical and mesh network. Other interexchange carriers tend to follow a hierarchical pattern, with extensive direct connections.

The interoffice portion of a CATV network does not involve switching at all (except in the sense that broadcast channels can be reallocated for different programming). Programming is normally broadcast from its source to a cable headend using satellites or other media. This broadcasting is augmented by sources of local programming that are directly coupled to the headend.

Today we see different topologies utilized in different segments of telecommunications and entertainment video networks, each with different capabilities. The customer premises component uses bus, tree, and ring structures. Network access utilizes a star topology for telecommunications and a tree topology for entertainment video. Finally, the interoffice portion of the telecommunications network uses a mixed hierarchical and mesh switching arrangement, while entertainment video uses a non-switched broadcast arrangement with a star topology.

Implications for IBN

What is the appropriate architecture for IBN? Let us start with several assumptions shared by IBN planners. First, in the early going IBN must coexist with existing LANs and other customer premises arrangements. Thus an IBN can be viewed as consisting of an inter-premises portion and a premises portion, with the premises portion being, in effect, a LAN, although one likely to be operating at speeds at or above 100 Mbps. The two IBN components will be interconnected through a gateway which, from a protocol point of view, may be simple or complex, depending on how similar the protocols are in the two components of the IBN. In an alternative picture the IBN would bring about the integration of customer premises and inter-premises networking so that, for instance, terminals currently connected to a LAN might in the future appear instead in a multi-drop configuration on an IBN access link. We will see below that there is another significant design consideration that fosters the division of the IBN into two components in this fashion.

Second, security and privacy considerations will mandate that customers will continue to have individual, physically separate access links. There will not be, for instance, a high-speed ring connecting customers together, to which they gain access on a shared-use basis as devices on a LAN do. Such sharing might lead to the illegal acquisition of transmitted information.

On the other hand, the use of such a ring for connecting the IBN equivalent of COs is not precluded, for the protection and isolation of individual customers is then in the hands of a disinterested third party, just as with today's public networks.

Third, the need to switch high-speed digital circuits with minimal delay mandates that the switching and routing functions be minimized, and done in a simple, distributed, and parallel fashion. The use of hierarchical topologies must be minimized or eliminated and transport will tend to be dispersed rather than centralized and hierarchical.

Given these assumptions, we can describe the following aspects of a proposed IBN architecture. First, there will be a premises portion which will appear much like those existing today. It will consist of either (or both) of the following. First, there may be one or more interconnected LANs with a connection to the external network through a high-speed network interface, or gateway. These LANs will be similar to those in use today but operate at higher speeds. A current issue is whether the LAN will utilize a ring or bus topology; either or both may exist with IBN. The recent adoption of a ring-based high-speed LAN standard known as Fiber Distributed Data Interface (FDDI) provides impetus to the ring topology. Second, there may be terminals of various types directly connected on an individual or multiple-tap basis to the network interface. It will be necessary to juxtapose an adapter between the terminal and the network interface to meet the IBN protocol specification.

A question arises about the configuration hypothesized above: should LANs continue to be isolated from the access portion of the network, with an interposed gateway? We believe so, for two reasons. By the time IBN is implemented there will be an enormous number of LANs and LAN-connected devices. Given that the LAN protocol is not suitable for extended distances and cannot support voice transport within metropolitan area networks, elimination of the gateway would necessitate a change in the LAN protocol and thereby require the changeover of all the devices connected to LANs. This is economically infeasible. In addition, it will be necessary to utilize distributed processing to the maximum degree possible. If the LAN gateway were eliminated, the processing involved in routing traffic between devices connected to the LAN would have to be performed in the IBN COs, thereby adding to the processing load, as well as increasing the transmission capacity required on the access links. Given the strong community of interest that led to the use of LANs in the first place, it seems likely that the best arrangement will be to continue to isolate the LAN traffic flow from the inter-LAN flow which must, of necessity, transit the metropolitan network.

A second aspect of the IBN architecture is that each premises will have an individual link to the serving IBN node. Because high data rates are required all the way to the premises, there may be less need for the multiplexing function currently appearing at the SAI, and its role will be limited to providing a cross-connection point to handle changes and growth over a period of time. Alternatively, the high-speed premises signals may still be multiplexed onto even higher speed feeder links. In order to provide protection against breakage of the access link, use of backup links to the same or another CO is likely to be widespread. In fact, in a configuration currently being employed, the links are carried on a fiber ring, providing two transmission paths between any two points.

Third, the IBN equivalent of the CO, referred to as an IBN node, will still route information from access links to interoffice links, and vice versa. It will, however, utilize a parallel processing approach rather than a centralized software-controlled routing function. This is discussed in greater detail below.

Fourth, there appear to be two alternatives for the metropolitan area portion of the interoffice network. First, the IBN nodes may be connected in a ring topology, operating in the fashion of an extended LAN ring. This is generally referred to as a Metropolitan Area Network (MAN) and is the subject of a current standardization effort by a subcommittee of the IEEE 802 Standards Committee. One or more gateways to inter-area networks would also appear on this ring. In this way, there would be no higher-order switching; rather, each node and gateway would pick off the information addressed to it. With a simple addressing scheme, this could be done largely via hardware registers, minimizing the processing involved.[10]

In the other alternative, the nodes would be connected in a star or hub arrangement, appearing superficially like today's tandem network. However, the hub would not switch in the fashion of today's digital switch, but would provide a cross-connect function between channels operating at a standard rate. The node would know the intended destination of the information being transmitted and would put this information in an appropriate time slot on its spoke to the hub so that the hub cross-connect would route the information to the destination node or gateway.[11] Again, parallel processing in the switch might be utilized to increase throughput.

Finally, beyond the metropolitan area, it is not yet clear whether the architecture and protocols of the MAN, or existing network capabilities will be utilized for long distance networks. Clearly, if there is a need for moving information between metropolitan areas at very high rates, some further evolution of these networks will be needed. On the other hand, there is a view that the need for high-speed transmission is of such a localized nature that existing network capabilities will suffice, augmented if necessary by the use of dedicated transmission pipes between MANs. In either case, there appears to be a natural division between metropolitan networks and long distance networks, brought about by the still-limited geographic range that the MAN protocol will allow.

In summary, the model IBN architecture for telecommunications will have many similarities to today's infrastructure and a few notable differences. The existing division between premises, metropolitan area, and long distance networks will remain; this division has a sound technical basis that transcends any regulatory or legal changes that may occur. The access architecture will still be a star topology for operational and security reasons. The notion of a LAN-like bus or ring topology, while appropriate to metropolitan area networking, may not extend significantly into long distance networking where a hierarchical structure is still appropriate. The main difference between IBN and today's network will reside in the metropolitan network, where a ring-like interconnection will replace the existing two-level switched hierarchical network.

If we extend this view to entertainment video, a few differences are notable. First, on the customer premises, all video devices will be connected to the network interface through an appropriate adapter, not via a LAN. In the access portion of the network, distribution will be handled either as part of the total stream of digital information — perhaps using a dedicated time slot — or in a separate signal carrying the information in analog form. Finally, the entertainment video portion of the IBN information flow will permit only gradual changes in cross-connection, while continuing to be received on a broadcast basis from the source of programming. It is likely to be handled by dedicated circuits which, in effect, bypass whatever switching mechanism is incorporated into IBN.

4. IBN Technology

Customer Premises Equipment

Three issues affect customer premises equipment as it pertains to IBN. The first has to do with the communications hardware and software in the devices attached to the premises portion of the IBN. These severely limit the input and output of data regardless of the throughput of the medium to which they are attached. Current emphasis is being placed on increasing the speed of the LAN media and little on the evolution of LANs to accommodate IBN. New developments in communications hardware and software will be necessary to take full advantage of the high data rates that these LANs and IBN can support.

Second, IBN is envisioned to integrate all forms of communications, including voice and low-speed data which use devices that today are not typically attached to LANs. As the existing effort to implement narrowband ISDN is making abundantly clear, the notion of adding a terminal adapter (TA) between existing terminals and the network is rather simplistic. This is so because of the expense involved. For instance, ISDN is being touted as a way to eliminate analog circuits and the associated need for modems; in fact, a modem, which often costs less than $100 and is built into a plug-in circuit board, would have to be replaced by an ISDN TA which may be significantly more expensive. While ISDN has the potential of supporting data communications at speeds much higher than those available today using voice-band analog data transmission, the added cost of doing so is dampening enthusiasm. The problem is compounded when one considers that IBN may require even more sophisticated devices to match telephones, low-speed data terminals, television, and other video devices to its wideband channels.

Third, another problem is evident today in attempting to connect LANs to other networks. Beyond simple bridging, the amount of processing required in the gateway can cause it to throttle throughput. This is particularly true when trying to match a connectionless network protocol to a connection-oriented protocol such as that being utilized in ISDN. In the former, normally used on LANs, individual packets are sent through a network without the prior establishment of a connection, or subsequent disconnection. In the latter, typical of circuit switched and X.25 connections, information is transferred only after a connection is established by the transfer of appropriate control packets. The contrasting protocols are quite different in the way they handle various events.

Given the need to match such protocol types, the gateway must perform call setup and participate in the error detection and correction process that is required by the connection-oriented protocol. It must also record and track the status of all the "virtual circuits" that exist during the time a call is connected. The processing involved will substantially reduce the throughput. The solution lies in picking an IBN protocol which facilitates LAN interconnection through the gateway. This is envisioned as a "datagram" or "connectionless" protocol, with data integrity being the responsibility of a higher-level protocol operating end to end in the communicating devices. A considerable amount of work remains to be done on such protocols if the potential throughput of IBN is to be fully realized.

Access and Interoffice Transmission

Unlike today's network, where there are significant differences between the technology of the access and interoffice portions, a single technology will be used for both portions of the IBN; that technology is fiber optics. There is simply no other transmission technology that can support the data rates envisioned for IBN. The only issue that remains is whether the transmission mode is to be multimode or single mode fiber. They differ in the way that light rays are propagated down the fiber. Single-mode fiber can support significantly higher data rates. Until recently it did not appear that there were sufficiently reliable and inexpensive light sources and detectors appropriate to the hostile access environment expected for IBN. Recent breakthroughs in commercially viable LED sources and detectors have mitigated this situation. This has led to the determination that single-mode transmission is the technology of choice, both in the access and interoffice portions of the network.

Commercial fiber optics transmission systems are already operating at 2.4 Gbps, with the maximum rate expected to increase to several gigabits per second within a few years. In the access network, the generally accepted view is that transmission rates of around 600 Mbps will be sufficient. This is based on an anticipated structure of four 150 Mbps channels for each premises. Three of these will be used to carry television or other video signals. The 150 Mbps rate is believed sufficient to carry digitally encoded HDTV. The other channel will be used to carry packetized information to and from the premises. With a single fiber, LAN interconnection at 150 Mbps will be possible. Alternatively, one can carry a multitude of digitized voice signals at this rate. This scenario assumes that entertainment video is carried digitally within IBN; we have already noted alternative solutions if this does not turn out to be feasible.

Switching

Switching at speeds fast enough to support transport rates in excess of 100 Mbps is more problematic. It is generally agreed that the switching must involve distributed and parallel processing. Attention is currently focused on a fast packet switching technique called Batcher-Banyan Routing.[12] In essence, this technique provides a compact, low-cost, fast, non-blocking packet switch implemented in VLSI. It does so by using very simple routing mechanisms in a multitude of parallel processors which do hardware switching based on the value of bits in an address. These bits are added to the front end of the packets being switched as they enter the switch. Experimental switches have already been developed that switch sixteen input ports to sixteen output ports at per-port speeds of 160 Mbps.[13]

Research is progressing on electro-optical switching that might replace VLSI-based technology. Such a system would have the advantage of low power consumption and high throughput. At present, obstacles to implementation are the size of the switching matrix, the complexity of control functions, and cost. The onset of commercially viable electro-optical switches is not projected to occur until well into the 1990s, especially given the current promise of VLSI-based technology.

Considerable effort is already being expended on the definition of protocols appropriate to IBN. Consensus is developing around a transmission format referred to by the name SONET.[14] It is a synchronous transmission protocol which will support a variety of speeds up to 2.5 Gbps, in integer multiples of the basic 50 Mbps rate. Of particular significance is the rate of 150 Mbps which accommodates the nominal channel rate envisioned by an IBN. SONET establishes a mechanism by which information to be transmitted is enveloped in control and signaling bits for passage through, and channel switching within, the network.

Within the SONET format, information and associated control overhead structures must be defined as well. Two different protocols are being considered at present, one called Asynchronous Transfer Mode (ATM) and the other, Synchronous Transfer Mode (STM). They differ in whether the structure of the fields within the SONET format are fixed or assigned dynamically according to the nature of the information to be transported.

The bottom line is that in all relevant technology areas — transmission, switching, protocols, and basic architecture — there is a significant amount of work being applied to the implementation of IBN. Work is progressing in the laboratories and in the field and includes associated national and international work on standards. The policy makers and strategists should assume that the elements of IBN will be commercially available within this decade.

5. Evolutionary Considerations

Let us present a scenario for the development of IBN based on the foregoing discussion of technology and architecture. First, it is assumed that premises LANs will continue to evolve toward higher speeds independently of other developments. Interconnection of LANs using high-speed gateways represents the second most likely IBN application, after video entertainment.

Secondly, regardless of the future of IBN, it is clear that fiber systems will be deployed in the access portion of both telecommunications and video entertainment networks. Given that fiber is being used to carry video signals, fiber systems become directly competitive with existing coaxial-based distribution systems. They require no realtime switching but only cross-connection of incoming television signals to access links. Even without any broadband switching capability there will be significant deployment of the access portion of an IBN, connected to both LEC COs and CATV headends. Furthermore, fiber is being deployed incrementally and as an overlay to the existing wire and coaxial distribution networks. A technique being utilized today employs fiber in the high cross-section feeder plant, with coaxial cable (for entertainment video) or wire pairs (for telecommunications) extending from the SAI to the individual premises. When fiber is put in place, electronics are installed which support the provision of voice and data circuits. These devices can be changed for higher-speed electronics when warranted by demand.

Operationally, the necessary planning, installation, and maintenance practices for fiber are being incorporated into existing LEC organizations. The degree to which these practices and associated operations support systems can accommodate IBN will depend on their similarity to the architecture and technology which is ultimately utilized.

Let us recall that the efficacy of utilizing digital transmission for television and other entertainment video has not yet been established. Many in the television industry believe that the distribution scheme will be analog due to the expense and inconvenience which would result from having to equip millions of existing sets with codecs. If this is true, there is a pending incompatibility with IBN plans for digital access systems.

Third, assuming that an appropriate IBN switching method becomes available, its deployment by the LECs is likely to be as an adjunct to existing switching systems in the same COs. Why would this occur? First, COs are located where the existing access network routes come together. Fiber is being deployed today to replace or augment the existing distribution media within the same physical facilities. Logically, then, the switching will be deployed in existing wire centers. Second, it will not be economical or technically feasible to replace the existing telecommunications network, at least in the short term. It is likely that there will always be a large number of premises that require only simple telephone services, perhaps augmented by low-speed data services such as alarms or telemetry. Unless the cost of IBN switching turns out to be inordinately low, those customers will best be served by existing narrowband circuit switching.

The overlay IBN nodes can readily be linked together by augmenting existing interoffice links with fiber. Customers requiring the switched capabilities of IBN can then be served by this overlay IBN network.

A similar statement may be made for cable television. Existing headends are the place where cable distribution trunks come together. If current entertainment video companies enter the IBN business they can be expected to install the switching systems in those headends.

We have already seen that the MAN protocol being developed for metropolitan areas may not be appropriate for long distance networks. Therefore gateways will be needed between the MANs and the long-haul transport networks. We do not anticipate any different organizational structures other than those which currently exist between the LEC and interexchange carriers. If other structures develop they will not be driven by technology considerations.

All of this suggests that IBN can be gracefully deployed, and in fact is already being deployed, as an overlay network, without disrupting existing organizational structures. Nevertheless, there are some organizational issues that arise from the foregoing scenario.

First, we might ask, should there continue to be a separation between the premises and access portions of the IBN as exists today? The apparent need to maintain a distinct on-premises network to keep switching as distributed as possible suggests that the answer is yes; or at least that there is no technical pressure to eliminate the separation. On the other hand, the growing similarity between the technology used in the LAN and the MAN suggests that synergy may result from having the same organization provide both functions.

Second, is the distinction between telephone and entertainment video companies appropriate? Most early deployment of fiber optics in the access network of the LECs is being justified in terms of its revenue potential as a medium for entertainment video distribution. The entertainment video industry has not been successful in carrying two-way data over coaxial cable, even after ten years of effort. Fiber optics appears to offer an opportunity for the merging of entertainment video and telecommunications, subject to the caveat, noted above, that the LECs and some members of the entertainment video industry are currently at odds on the issue of digital versus analog. Will both types of companies be needed in the future in light of this technological merger?

Third, should the LECs have the exclusive right to deploy a new access technology? One might hold that to avoid the same local access bottleneck that has always existed in the provision of telephone service, alternative suppliers of local IBN services should have the opportunity to deploy the fiber systems. One way in which this might occur naturally is for the current providers of LANs to start expanding their networks outside the premises boundary. There is an issue brewing in many state jurisdiction over what constitutes a "premises," particularly in connection with the provision of so-called "shared tenant services". Another possibility is that entertainment video companies, not LECs, will deploy a fiber access network.

Fourth, one might ask why the LECs should be the exclusive providers of IBN metropolitan area networks. Given that the existing switching systems are not adequate to support IBN and must be augmented by IBN nodes, should not the provision of those nodes be open to competition? This could be facilitated by allowing colocation of other entities within LEC COs and by the continued sharing of pole and duct facilities. Even without colocation — which the LECs are firmly on the record as opposing — the relative distance-insensitivity of fiber system performance and cost makes it possible to imagine a competitive situation. Other providers could locate next to a LEC CO, where significant IBN demand exists, and connect to and through the CO to other locations. Or they could operate fiber rings around data-intensive portions of a metropolitan area and bypass the LEC COs altogether.

Organizationally, IBN is different enough in both the transmission and switching realms and has such minimal deployment today that one can envision competitive access. Rather than assuming that an overlay network falls under the exclusive province of the LECs, it seems equally reasonable to assume that other firms such as entertainment video companies will provide competition. At the same time, separate IBN providers are not necessarily appropriate. Nor should it be assumed that the providers of the premises portion of the IBN should be separate from the inter-premises provider.

6. Conclusions

We have defined an IBN as a network supporting data applications involving the transport of digital information, and the distribution of television and other video information at very high rates. These rates are currently viewed to be on the order of 600 Mbps in the access network and probably in the gigabits per second range for interoffice transport. The transmission technology is clearly fiber optics. The switching technology has not been developed although a great deal of work is focused on the notion of fast packet switching, a generic name for several embryonic technologies such as Batcher Banyon routing. Experimental models are showing promising results. In parallel with the effort to define, design, and implement public IBNs, developments in the LAN arena are keeping pace, insuring that a combined IBN will be feasible in the 1990s time frame.

IBN raises a number of organizational issues which spill over into the policy arena. It is not clear that there will continue to be a need for separate telephone and entertainment video networks. If not, what is the appropriate organizational arrangement for the sharing of facilities and the joint provision of services? It is unclear why the telephone companies should have exclusive right to deploy IBN, given its differences from current networks.

Notes

1. As a convenience, in differentiating between the technology and architectures utilized by today's local exchange and interexchange carriers from those of CATV operators, we will refer only to the former as telecommunications; even while recognizing the frequent practice of including both types of service under that designation. We will use the term "entertainment video" interchangeably with CATV.

2. See John McQuillan, "What's Wrong with ISDN?," *Business Communications Review*, January/February, 1987.

3. Alex Best, "Cable Television and Fiber Optics: Has Fiber's Day Finally Arrived?," talk presented at the Eastern Communications Forum, Rye Brook, New York, May 4, 1988.

4. Christine Hemrick, et al., "Switched Multi-megabit Data Service and Early Availability Via MAN Technology," *IEEE Communications Magazine*, April, 1988, p. 9.

5. W. Stallings, *Local Networks: An Introduction*, MacMillan Publishing Company, New York, N.Y., 1987.

6. "Notes on the BOC Intra-LATA Network - 19896," Bell Communications Research,Inc., Technical Reference TR-NPL-000275, Chapter 12, 1986.

7. Marvin A. Sirbu and David P. Reed, "An Engineering and Policy Analysis of Introducing Fiber into the Residential Subscriber Loop," in this volume.

8. M. Frame, "Broadband Access," talk presented at the Eastern Communications Forum, Rye Brook, N.Y., May 4, 1988.

9. G. Ash and V. Mummert, "AT&T carves new routes in its Nationwide Network," *AT&T Bell Laboratories Record*, August, 1984, p. 19.

10. R. Newman et al., "The QPSX MAN," *IEEE Communications Magazine*, April, 1988, p. 20.

11. S. Cheng and L. Wu, "Reconfigure in a Flash!," *Telephony*, May 18, 1987, p. 76. Adaption of a presentation at the International Telecommunications Symposium, Taipei, Taiwan, September, 1986.

12. S. Cheng and L. Wu, ibid., pp. 79-80.

13. A. Vaidya, "Wideband Packet Switching," talk presented at the Eastern Communications Forum, Rye Brook, N.Y., May 3, 1988.

14. L. Campbell and C. Engineers, "Synchronous Transmission Beyond the 1990s," *Lightwave Magazine*, December, 1987, p. 44.