New Directions in Communications (or Which Way to the Information Age?)

Jonathan S. Turner

Information transport networks should provide a generally useful set of communications capabilities that can be used as a foundation on which to build a large class of different applications. This article discusses the need for these networks and considers ways of implementing them

The Information Age had been characterized as a time of exploding demand for communications applications of all kinds. These new demands create both great opportunities and difficult challenges for suppliers of telecommunications services and equipment. The challenges are difficult because our current communications technologies have been designed around specific applications and are difficult to adapt to new ones. Even when an existing system can be adapted to a new application, the required investment may preclude development unless there is a large and clear demand. The current practice of responding to diverse needs by developing application-specific communications services is unworkable in an environment of rapidly changing demands. What's needed are networks that provide basic communications capabilities in an application-independent fashion. Such *information transport networks* should provide a generally useful set of communications capabilities that can be used as a foundation on which to build a large class of different applications. In this article, we discuss the need for such networks and consider some possible ways to implement them.

The communications networks that are most widely used by the general public today are the telephone and cable television networks. One of the first things one notices about these networks is that each is oriented toward a particular application and while they perform their assigned function well, they are poorly suited to anything else. Even packet-switched data networks, while relatively flexible, have such a limited performance range that they can be applied to only a rather narrow class of applications.

The predominant reaction currently to the limitations of individual *applications networks* is to deploy several of them in parallel. Thus, we currently have telephone and CATV networks operating in parallel and we are starting to see the deployment of data networks for use by the general public. This "solution" is inefficient and expensive. What's worse, it's unlikely to satisfy long term needs, since new applications are inevitable and the existing and planned networks are unlikely to be able to accommodate them. What's needed is an integrated network capable of supporting a variety of applications, not some haphazard collection of parallel and mutually incompatible applications networks.

The need for integrated communications systems has been recognized for some time. Most attention currently centers on Integrated Services Digital Network or ISDN [2,5,13] Work on ISDN has resulted in a standard for a digital line interface, comprising two 64 Kb/s circuit switched channels for voice and bulk data and a 16 Kb/s packet switched channel for data and control information. The trouble with ISDN is that it is not really an integrated network. The only thing integrated

---

This research was supported by Bell Communications Research and Italtel SIT. An early version of this paper appeared in the Proceedings of the Zurich Seminar on Digital Communications, 1986.

0163-6804/86/1000-0008 $01.00 © 1986 IEEE

out it is the physical transmission path from the
ıtral office to the customer's premises and even here,
parate channels are prescribed. The communications
stems required to support this interface must provide
parate switching mechanisms for circuit switched
d packet switched channels. The only integration
ssible is at the level of physical packaging. Little
ought has been given to providing higher rate
vices such as high speed data applications or video.
ıe suggestion that these services could be handled by
provision of a large number of 64 Kb/s circuit
itched channels is unrealistic.

The next generation of large scale public communi-
tions systems should be more than a repackaging of
current generation. If we ever expect to develop
stems that can accommodate new applications easily,
must stop designing each new network to a partic-
ır application and focus instead on networks that
ovide general capabilities and offer a wide range of
rformance options. Such an information transport
twork should be able to provide connections in a
de range of bandwidths. Applications, such as telem-
y require just a few bits per second, while video may
quire 100 Mb/s or more. An information transport
twork should be able to handle such extremes, as
ll as everything in between. If such a network is to
commodate broadcast video, it must provide broad-
st, as well as point-to-point connections. A broadcast
nnection can be viewed as a special case of a more
ıeral multi-point connection, joining an arbitrary
mber of users. Such a multi-point connection could
used for voice or video tele-conferencing, in addition
broadcast. A network of this sort could handle all
applications currently provided using multiple
plications networks. More important, it would be
xible enough to allow easy introduction of new
plications as they come along.

ew Switching Technologies

What are some of the options available to would-be
igners of next-generation communications systems?
[12], Kulzer and Montgomery discuss several possi-
ities, placing them in the spectrum shown in Fig. 1.
chniques toward the right end of the spectrum
ovide increasing flexibility to handle variable rate
ormation, but require more processing. The region
ır the center of the spectrum is labelled statistical
tching and contains technologies that can transport
sty information streams without the full function-
y of conventional packet switching.

To those with a background in telephony, Multi-
e Circuit Switching (MRCS) seems a natural choice
new communications systems. MRCS provides
nections having bandwidths equal to an integer
ltiple of some basic rate, such as 8 or 64 Kb/s. The
r specifies a transmission speed when the call is set-
and the network provides enough channels to
sfy the request. As Kulzer and Montgomery point
, there are several problems with MRCS. First is the
ice of the basic rate; many services require a low rate
h as 1 Kb/s, but this can imply a long delay due to
large frame size required for time-division multi-
plexing. It also implies a large overhead for establishing high speed connections, since these must be implemented using multiple channels, each of which must be set-up individually. Even with a fairly large basic rate of 1 Mb/s, a video connection might require the establishment of 100 channels. MRCS is also ill-suited to applications with bursty transmission characteristics. Applications such as remote file access require occasional transfer of bursts of data at high rates such as 10 Mb/s. Dedicating a high speed connection to such applications is costly and inefficient. Using a lower speed channel yields efficiency, but only by sacrificing performance.

If we take the next step to the right on the spectrum in Fig. 1, we come to Fast Circuit Switching (FCS). As with MRCS, users request connections having bandwidth equal to some integer multiple of a basic rate. In FCS however, the system does not allocate the required channels until the user has some information to send. Thus, FCS allocates bandwidth dynamically among a group of users, allowing efficient sharing of the transmission facilities. Of course, there may be occasional peak traffic periods when the network cannot satisfy all users' requests. When this happens, one or more requests are denied. This kind of switching has been termed "burst switching" by Amstutz [1] and Haselton [6].

Fast Packet Switching (FPS) is the next option on the spectrum. As in conventional packet switching, FPS uses the transmission facility as a "digital pipe," which carries short packets of information one after another. Information in the header of each packet identifies which of many logical connections the packet belongs to. With this multiplexing scheme, connections of arbitrary bandwidth are accommodated in a simple and natural way. A key aspect of FPS is the recognition that the high speed and low error rate of modern digital transmission facilities allow simplification of the communications protocols used in conventional packet switches. These simplifications make possible the construction of hardware protocol processors. High speed transmission facilities also dramatically reduce the queueing delays inherent in packet switching. Another key element is the observation that high speed computer interconnection networks originally designed for large parallel computer systems [3], are ideally suited to large high performance packet switching systems. FPS has been developed by a group at Bell Laboratories and is described in references [8,10,12,16, 19,20,27]. Similar work had been done by groups at Lincoln Laboratories and Bolt Beranek and Newman [18].

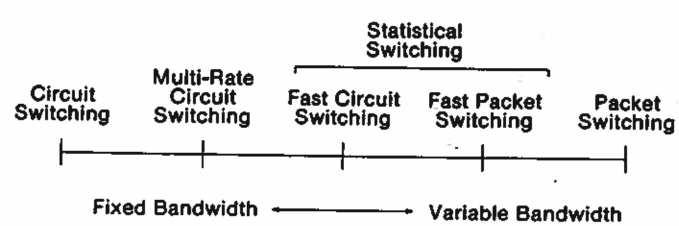

Fig. 1. Switching Technology Spectrum.

Multi-point Networks

While the technologies described in the previous section can serve a wide class of applications, they lack the ability to handle broadcast or multi-point connections. In this section we consider some of the problems inherent in the design of multi-point networks. To make the discussion concrete, we put this in the context of an extension of the Fast Packet Switching technology, which we call Broadcast Packet Switching (BPS).

Video based applications pose two major challenges. First, how to provide the much larger bandwidths required and switch them effectively. Second, how to provide for broadcast, or more generally multi-point connections in addition to the conventional point-to-point connections. While the switching of high bandwidth channels poses a challenging engineering problem, its solution requires no real conceptual departure from the earlier work. The problem of multipoint connections is a more difficult and interesting one, which raises a variety of new issues. These issues can be placed into three broad classes:

- *Architecture of switching systems*—In point-to-point communications networks, a switching system must be able to connect any incoming channel to any outgoing channel. To support multi-point connections, a switching system must be able to connect any incoming channel to any subset of its outgoing channels.
- *Connection management*—In point-to-point networks the algorithms that manage the establishment of connections have a fairly simple data management problem. To handle multi-point connections, these algorithms must cope with connections that branch in many different directions. Furthermore, they must provide a connection mechanism that is flexible enough to support point-to-point, broadcast and conference connections as special cases.
- *Network control issues*—Many of the global network control problems change radically in networks that provide multi-point connections. For example, routing in point-to-point networks can be treated as a shortest path problem. In networks providing multi-point connections, what is needed is a method of finding the shortest tree connecting a given set of endpoints.

Broadcast packet switching is a new switching technology that can be used to implement multi-point networks. The major components of a broadcast packet network (BPN) are identified in Fig. 2. Access to the network is provided through Network Interfaces (NI), which provide concentration, network protection and accounting functions. The switching function is provided by Packet Switches (PS), which are configurable over a wide range of sizes. Large systems should be able to support 50,000–100,000 users. Transmission is over Fiber Optic Links (FOL) on which a large number of logical channels are statistically multiplexed. The network provides several basic communications services:

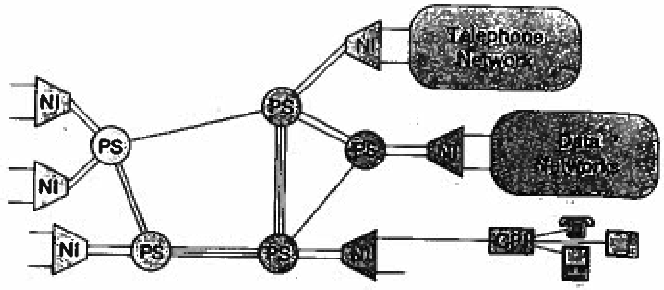

Fig. 2. Network Architecture.

- *Point-to-point connections*—These are two-way connections joining pairs of users. The user requests a channel capable of providing a certain average bandwidth and the network allocates the requested bandwidth to the connection. If the requested bandwidth is not available, the connection is blocked.
- *Datagrams*—These are individual packets, not associated with a pre-established connection. The network makes an effort to deliver them, but does not guarantee delivery.
- *Broadcast connections*—Any user can set-up a broadcast source that other users can then connect to. The average bandwidth of the source must be specified when it is established, and can range from a few bits per second up to the speed of the network's transmission links. There is no limit on the number of users that can receive a given broadcast signal. Thus, it is suitable for a variety of applications including commercial television distribution.
- *Conference connections*—A conference connection can be viewed as a multi-way broadcast. Packets sent by any participant in the connection are received by all others. They can be used for voice or video teleconferencing.

Figure 3 illustrates how broadcast (and conference) connections are implemented. Two broadcast sources are shown at the top of the figure. At various points in the network, the signals are split into multiple copies, which are ultimately delivered to the appropriate users. Thus, each broadcast source and its associated endpoints induce a tree in the network. When a user disconnects from a broadcast channel, the corresponding branch of the tree is removed. When a user requests connection to a channel, the network attempts to find a nearby branch of the broadcast tree and connects the user at the nearest point. The broadcast trees grow and shrink as usage patterns change, but for widely distributed channels, one can expect the bulk of the activity to be concentrated near the leaves.

Simple link level protocols are used to facilitate hardware implementations of basic switching and protocol functions. In particular, the link level protocols exclude error correction and flow control, eliminating the need for synchronization of state information between the two endpoints of each link. Large, fixed-length packets are also used to further simplify implementation of the protocol and switching hardware.

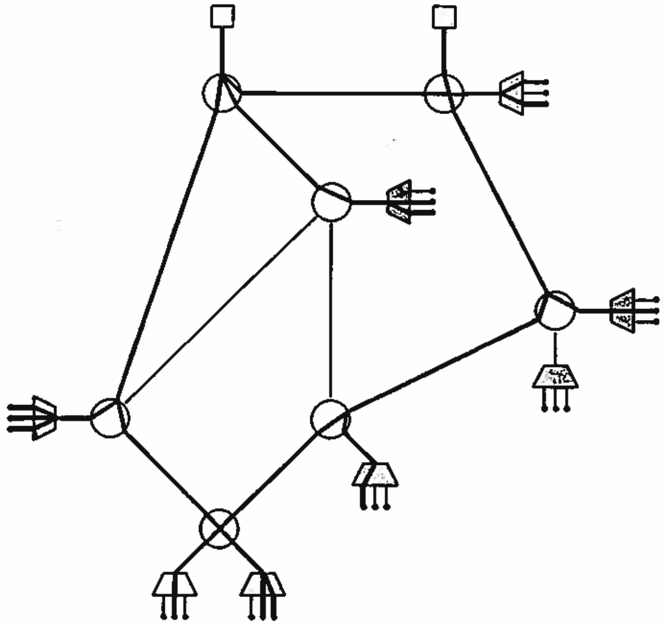

Fig. 3. Broadcast Trees.

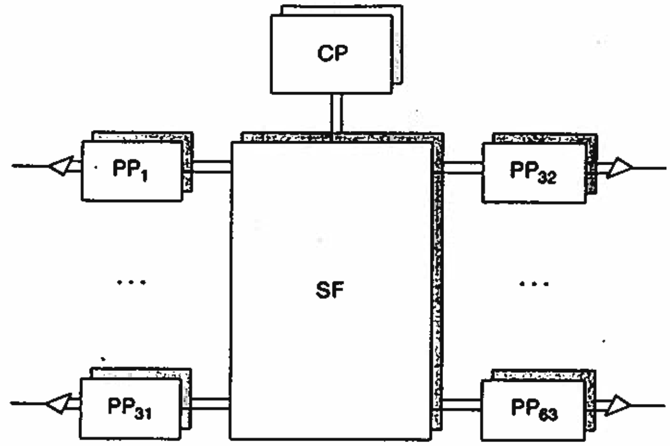

Fig. 5. Switch Module.

Architecture of Switching Systems

We have been studying a class of modular architectures for large packet switching systems. The attraction of these architectures is that they permit easy growth of a switching system over a wide range of sizes. An example of one such architecture is shown in Fig. 4. It contains a number of identical components called Switch Modules (SM) which provide the basic switching functions. The SM's are of two types: Front-end Switch Modules (FSM) and Back-end Switch Modules (BSM). The FSM's provide all the external connections while the BSM's interconnect the FSM's. The number of SM's can vary over a wide range, allowing a small intial configuration to grow as demand increases. The Cross Connect (XCON) is a special purpose space division switching element that allows reconfiguration to be done without manual recabling. Reference [30] describes a simple design for the XCON that has linear growth complexity rather than the quadratic complexity one would expect for a crossbar switch. This somewhat surprising design is made possible by the limited function required of the XCON. The Supervisory Processor (SP) provides office administration and craft interface functions.

To permit the construction of large packet switching systems in this modular fashion, the SM's must have the ability to establish connections and to process individual packets. Each SM is largely independent of the others in the system. The interface between SM's in the same switching system would be the same as the interface between different switching systems. This approach permits a variety of switching system architectures, with the SM acting as the basic building block.

The design of the SM is shown in Fig. 5. The SM terminates 63 fiber optic communications links (FOL), each operating at a speed of 100 Mb/s and engineered for an occupancy of 80 percent giving a raw throughput of approximately 5 Gb/s. The Packet Processors (PP) perform the link level protocol functions, including the determination of how each packet is routed. The Switching Fabric (SF) is the heart of the SM. It has the ability to route point-to-point packets to the proper outgoing FOL and the ability to replicate packets belonging to braodcast connections and route each copy to the appropriate FOL. Connection Processor (CP), is responsible for establishing connections, including both point-to-point and broadcast connections. To do this, it exchanges control packets with CP's in neighboring SM's and controls the actions of the PP's and SF by writing information in their internal control

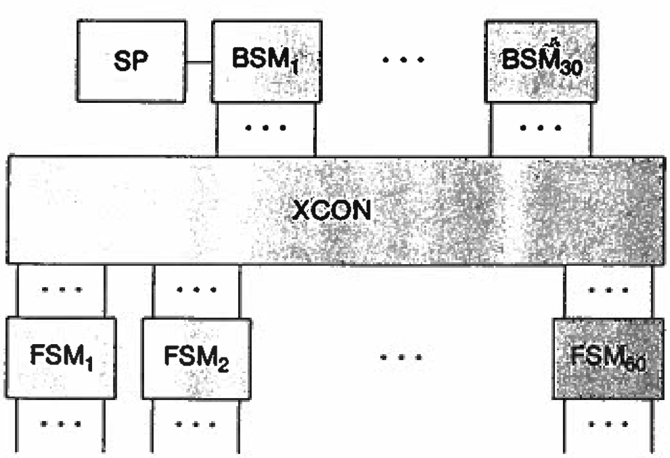

Fig. 4. Small Packet Switch.

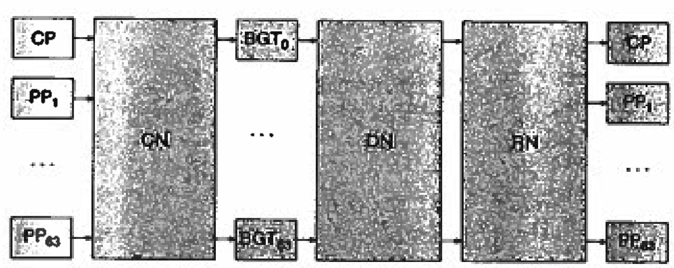

Fig. 6. Switch Fabric.

tables. It also performs a variety of administrative and maintenance functions.

Each packet received by a PP from an FOL contains a Logical Channel Number (LCN) identifying the connection it belongs to. Upon reception, several additional header fields are added to the packet including a routing field, which contains information needed to route the packet through the switch fabric. For point-to-point packets the routing field would contain the number of the outgoing FOL and a new LCN value. For broadcast packets it would include the *Fanout*, that is the number of copies required and an identifier called a Broadcast Channel Number (BCN).

A block diagram of the Switch Fabric (SF) is given in Fig. 6. It contains four major components, a *Copy Network*, a set of *Broadcast and Group Translators*, a *Distribution Network* and a *Routing Network*. When a broadcast packet having a fanout of k passes through the Copy Network (CN), it is replicated so that k copies of that packet emerge from the CN. Point-to-point packets pass through the CN without change. The Broadcast and Group Translators (BGT) modify each broadcast packet by mapping the packet's BCN to an outgoing FOL number and outgoing LCN. This new information is inserted in the packet's routing field. The Distribution and Routing Networks (DN,RN) use the outgoing FOL number to route packets to the proper outgoing FOL. The RN does the actual routing, while the DN distributes traffic evenly across the RN to prevent internal congestion. Each of the networks comprises an $O(n \log n)$ complexity network with buffered switching elements such as a delta or banyan network [3]. Although identical in topology, they differ in how they process packets. The RN routes packets in the conventional fashion, the DN distributes them randomly and the CN performs replication. A detailed description appears in [27].

Connection Management

Connection management refers to the collection of algorithms used to create and maintain connections among users. Point-to-point, broadcast and conference connections can all be viewed as special cases of a more general multi-point connection concept. The connection management algorithms must be capable of handling connections with an arbitrary number of endpoints (possibly millions) and must be able to respond rapidly to changes in the configuration of a connection. These considerations imply the need for distributed control algorithms that allow each switching system to manage its part of a connection with only local information. Multi-point connections must be able to change in complex ways. For example, it must be possible to add new users to a teleconference connection. Similarly, it must be possible to reconfigure a broadcast video connection in response to changing demand. Note that this second case differs in that a request to join such a broadcast connection must come from a user who is initially outside the connection.

Of the three types of connections considered above, the conference connection is the most general. In this type of connection, packets sent by any participant are forwarded by the network to all others. By adding suitable control mechanisms, this type of connection can be tailored to other uses. One addition is to allow users in a connection to send private messages. This can be implemented by having the sending user include a destination address on private messages, which the network uses to restrict delivery to the specified endpoint. To obtain more flexibility, we introduce the notion of subchannels within a connection. A subchannel is characterized by its average bandwidth, the set of endpoints that can transmit on the subchannel and the set that can receive. To implement a broadcast connection, we use one subchannel on which the only transmitter is the broadcast source, while all other endpoints can only receive. To add a low bandwidth upstream capability, we provide a second subchannel on which the broadcast source receives and all the others transmit. The concept of a connection containing multiple subchannels is very flexible and allows bandwidth to be allocated in the network only where it is needed.

It's useful to think of the network and the connections it provides at several levels of abstraction. The Level N abstraction defines the services that the network offers the user and the protocol used to provide them. At this level we can view the network as a collection of terminals interconnected by a single central node in a star configuration. The connection at this level is characterized by its endpoints and its subchannels.

At Level $N-1$, we view the network as a graph in which the internal nodes correspond to switching systems and the edges correspond to FOL's or groups of FOL's interconnecting the nodes. At this level a connection is characterized by its topology in addition to its endpoints and subchannels. The protocol used by switching systems to establish and manage connections is defined at this level.

At Level $N-2$, we decompose the individual nodes into subsystems. Since different switching systems will exhibit a varitey of internal structures, the mechanisms used to implement connections will vary.

It's also useful to consider Level $N+1$, which includes application-specific concepts. Here for example, is where the protocols for voice communications would be defined. Different protocols would be defined for data and video comunication. It is our feeling that this level should lie outside the network proper. All application-specific information and protocols would reside in the users' equipment, with the network providing the Level N abstraction used by the higher levels to implement their services.

There are several other sets of issues that arise in large-scale public networks. These are discussed briefly below.

- *Recovery*—In the telephone network, equipment failures can cause connections to be broken off by the network. Broadcast networks are likely to require a lower rate of cutoff connections, since the connections can last for long periods of time and involve a large number of users. It may be necessary to implement strategies that re-establish broken connections when a packet switch fails.

- *Naming*—In a broadcast network, there must be mechanisms for identifying connections so that users can add onto connections from "outside". This also means that the network must provide a mechanism by which users can discover connection identifiers. It may also be useful to have the network assign names to users, rather than just to network terminals. This facilitates more flexible routing mechanisms and is useful for authorization.
- *Authorization*—The possibility of connections that can be joined by outside users raises a variety of authorization questions. What mechanisms does the network use to decide whether or not to permit a requested connection? Some applications will forbid connection by outsiders, others may allow them without questions and still others may require explicit authorization for each new user. Identification of users is an important part of authorization. Should the network identify new users to those on the connection and if so, how is that identification accomplished?
- *Security*—Should the network attempt to protect users' data from spying or tampering or should that function be provided outside the network proper? If the network does not protect user data, how does it protect its own data? Should control messages between users and the network be protected?

Network Control Issues

Network control refers to global issues involved in controlling a large communications network. These problems are typically handled by a combination of hardware and software mechanisms. We focus here on two network control problems, routing of multi-point connections and congestion control.

Routing of Multi-Point Connections

The problem of routing connections in broadcast networks is quite different from routing in point-to-point networks. In this section, we first review the routing problem in point-to-point networks with uniform connections, then consider the complications introduced by adding variable bandwidth and multiple endpoints.

Consider a point-to-point network in which all connections require the same amount of bandwidth. Such a network can be described formally as a graph in which each edge has both a capacity and a length. A set of connections for such a network is simply a collection of vertex pairs. A feasible route assignment is an assignment of each connection to a path in the network joining the connection's endpoints that doesn't exceed the capacity of any edge. That is, the number of connections using any particular edge must be bounded by the capacity of that edge. An optimum routing algorithm is one that can find a feasible assignment whenever one exists. Of course, this version of the problem is a static one. In a real communications network, the set of connections changes with time and the network must implement a routing policy that manages the changing set of connections in a way that makes it unlikely that a new connection will be blocked. In the interests of efficiency, it is generally assumed that once a connection has been assigned a route, that assignment will remain fixed as long as the connection is present. These considerations lead to a routing policy based on the heuristic strategy of routing connections by the shortest path available at the time the connection is established.

If connections can have an arbitrary bandwidth associated with them, the routing problem becomes a bit more complicated. One must now consider the network to be a graph in which vertices can be joined by multiple edges. To prevent blocking of connections with large bandwidth requirements, new connections should be assigned to the fullest edges with sufficient capacity along the assigned route. This strategy preserves large blocks of bandwidth for use by high speed connections. (Note the similarity between this problem and the problem of memory allocation in computer systems.)

In broadcast networks, a connection can involve an arbitrary number of endpoints. A feasible route assignment for a set of connections is an assignment of each connection to a subtee connecting its endpoints, in a way that does not exceed the capacity of any edge. As in the case of point-to-point networks, connections come and go over time, and so the appropriate routing policy is to assign each connection to the subtree with shortest total length available at the time the connection is established. This can be viewed as a generalization of the Steiner tree problem in graphs [4]. This problem is known to be NP-complete, meaning that there is unlikely to be an efficient algorithm that can always find an optimal solution. On the other hand, there is at least one efficient algorithm that yields solutions that are close to optimal [11].

Connections in broadcast networks are dynamic in another way. They grow and shrink with time as individual endpoints come and go. The challenge is to maintain a good connection topology without doing a great deal of recomputation each time an endpoint is added or dropped. In addition, a practical algorithm must be possible to implement in a distributed fashion, with each node making decisions based on local information. The simplest strategy is to add new endpoints by joining them to the connection by the shortest available path and dropping branches of the connection tree when endpoints drop out. While this handles the dynamic nature of the problem and is suitable for distributed implementation, it can perform poorly in the worst case. Its performance in typical cases needs to be evaluated.

There are other considerations that affect routing algorithms for broadcast networks. The use of subchannels with differing bandwidth requirements in a connection is one. Another is the need to limit the diameter of a connection (that is, the maximum distance separating two endpoints) to prevent excessive delay.

Congestion Control

A principal advantage of packet switched networks is their ability to dynamically allocate bandwidth to the users who need it at a particular instant. Since networks are subject to rapid statistical variations in demand, care must be taken to ensure acceptable performance under conditions of peak loading. The problem of controlling the effects of peak loading is particularly severe in connectionless networks, which have only a limited ability to restrict total demand. In a connection-oriented network, bandwidth can be allocated to new users as connections are established and new connections can be refused if there is insufficient bandwidth available, thus ensuring predictable performance, once a connection is set up. Connectionless networks, on the other hand, respond to overloads by degrading the performance seen by all users. Congestion control refers to the collection of methods used to ensure each user acceptable performance under a variety of load conditions. The high speed and multi-point connection capability of broadcast packet networks place new demands on congestion control methods.

We believe that an effective congestion control system requires several specific methods, each acting on a different time scale. Long term overloads are prevented by the allocation of bandwidth to connections and the refusal of new connections unless the needed bandwidth is available. This means that the network must provide a mechanism for users to specify their bandwidth needs and an indication of the burstiness of their transmissions, and must enforce limits to prevent users from exceeding their allocations. One way to do this is discussed below. Short term demand variations are handled by buffering within the network. Assuming the architecture described in the previous section, a FOL buffer with room for 32 packets can fill up in about 3 ms. A complete congestion control system must also provide mechanisms for handling peak periods of intermediate duration.

A key part of bandwidth allocation is the mechanism used to specify the needed bandwidth and limit users to their allocations. Perhaps the simplest approach is the so-called "leaky bucket" method. A counter associated with each user transmitting on a connection is incremented whenever the user sends a packet and is decremented periodically. If the counter exceeds a threshold upon being incremented, the network discards the packet. The user specifies the rate at which the counter is decremented (this determines the average bandwidth) and the value of the threshold (a measure of burstiness). These two numbers can be used by the network to allocate bandwidth and limit the flow of packets into the network.

Packet priorities offer one promising method of coping with peak periods lasting up to a few seconds. During high demand periods, the network can preferentially discard low priority packets. If half the packets on a specific FOL have low priority, that FOL can tolerate peak periods of arbitrary duration without losing any high priority packets. Priorities can be used to advantage for signals containing large amounts of redundant information. For example, video signals can be transmitted with the high order bits of each pixel carried in high priority packets and the low order bits carried in low priority packets. Occasional loss of low priority packets would probably be imperceptible. Periods of a few seconds during which many low priority packets are lost, are likely to be perceptible, but only mildly annoying to the viewer. Similar methods have been used effectively for packet transmission of voice signals.

Longer term control can be obtained by having the network inform users of peak loading conditions and request that they reduce their rate of transmission. Ideally, users should be allowed to continue the same average rate of transmission, but would have to reduce the burstiness of their signals. For example, if the leaky bucket method were used to limit the flow of packets into the network, the threshold at which packets are discarded could be reduced in response to heavy loads.

In point-to-point networks, the flow of packets into the network can be controlled entirely at its edges. In networks with multi-point connections, this doesn't appear to work, since packets from many users in a single connection can converge onto a single link and exceed the allocated bandwidth at that point. The brute force solution to this problem is to measure and limit the bandwidth of each connection at each link in the network. Another solution is to measure and limit the combined bandwidth used by packets entering the connection and those leaving it at each access point. If the leaky bucket method is used, a user's counter is incremented every time he either sends a packet or receives one. This has the effect of limiting the total bandwidth on the connection. This approach can be extended to multi-channel connections by having separate counters for each channel.

Summary

The current proliferation of multiple application-oriented networks is inefficient, expensive and unlikely to satisfy long term requirements. As the need for new applications grows, the limitations of current systems will become increasingly troublesome. The Integrated Services Digital Network is at best a stop-gap measure that will postpone the problems for a few years. An effective solution requires the development of flexible information transport networks, capable of providing connections of arbitrary bandwidth and with multiple endpoints. We have described a new switching technology that can be used to implement such a network and have discussed some of the research issues raised by this technology.

References

[1] Stanford R. Amstutz, "Burst switching—an introduction," *IEEE Communications Magazine*, pp. 36-42, Nov. 1983.

[2] Maurizio Dècina, "Progress towards user access arrangements in integrated services digital networks," *IEEE Trans. on Comm.* vol. COM-30, no. 9, 2117-2130, Sept. 1982.

[3] Tse-yun Feng, "A survey of interconnection networks," *Computer*, vol. 14, no. 12, pp. 12-27., Dec. 1981.

[4] Michael R. Garey, and David S. Johnson. *Computers and Intractability, a Guide to the Theory of NP-Completeness*, W. H. Freeman, 1979.

[5] "ISDN user-network interfaces," *Proc. of the International Conf. on Comm.*, pp. 1072-1075, June 1985.

. Fletcher Haselton, "A PCM frame switching concept ading to burst switching network architecture," *IEEE ommunications Magazine*, pp. 13-19, Sept. 1983.
urt Haserodt and Jonathan Turner, "Connection man- ;ement in broadcast packet networks," Washington niversity, Computer Science Department, technical re- ɔrt WUCS-85-9, Aug. 1985.
'illiam L. Hoberecht, "A layered network protocol for ıcket voice and data integration," *IEEE Jour. on ·lected Areas in Comm.*, vol. SAC-1, no. 6, pp. 1006- ı13, Dec. 1983.
lan Huang and Scott Knauer, "Starlite: a wideband gital switch," *Proc. of Globecom '84*, pp. 121-125, Dec. '84.
.h-Chyun Jenq, "Performance analysis of a packet ɾitch based on a single-buffered banyan network," *'EE Jour. on Selected Areas in Comm.*, vol. SAC-1, no. pp. 1014-1021, Dec. 1983.
Kou, G. Markowsky, and L. Berman, "A fast algo- hm for Steiner trees," *Acta Informatica*, vol. 15, pp. 1-145, 1981.
hn J. Kulzer, and Warren A. Montgomery, "Statistical itching architectures for future services," *Proc. of the ternational Switching Symposium*, May 1984.
C. Luetchford, "CCITT recommendations—network ɔects of the ISDN," *Proc. of the International Conf. Comm.* pp. 1067-1071, June 1985.
:k F. McDonald, Edwin H. Rogers, Kenneth Rose, d Andrew J. Steckl, "The trials of wafer-scale inte- ıtion," *IEEE Spectrum*, pp. 32-39, Oct. 1984.
metrios Michalopoulos, 'Second-generation cray com- ter unveiled," *Computer*, p. 116, Aug. 1985.
ırren A. Montgomery, "Techniques for packet voice ıchronization," *IEEE Jour. on Selected Areas in mm.*, vol. SAC-1, no. 6, pp. 1022-1028, Dec. 1983.
Oktay, and H. C. Kammerer. "A conduction-cooled ›dule for high performance LSI devices," *IBM Jour. Research and Development*, pp. 55-66, Jan. 1982.
Rettberg, C. Wyman, D. Hunt, M. Hoffman, P. rvey, B. Hyde, W. Clark, and M. Kraley, "Develop- nt of a voice funnel system: design report," Bolt anek and Newman, Report no. 4098, Aug. 1979.
ıathan S. Turner, and Leonard F. Wyatt, "A packet work architecture for integrated services," *Proc. of ›becom '83*, pp. 45-50, Nov. 1983.
ıathan S. Turner, "Packet load monitoring by trunk trollers," United States Patent #4,484,326, Nov. 20, 4.
athan S. Turner, "Packet switching loop-around ɾork and facilities testing," United States Patent :86,877, Dec. 4, 1984.
athan S. Turner, "End-to-end information memory arrangement in a line controller," United States Patent #4,488,288, Dec. 11, 1984.
[23] Jonathan S. Turner, "Interface facility for a packet switching system," United States Patent #4,488,289, Dec. 11, 1984.
[24] Jonathan S. Turner, "Packet error rate measurements by distributed controllers," Unites States Patent #4,490,817, Dec. 25, 1984.
[25] Jonathan S. Turner, "Fast packet switch," United States Patent #4,491,945, Jan. 1, 1985.
[26] Jonathan S. Turner, "Fast packet switching system," United States Patent #4,494,230, Jan. 15, 1985.
[27] Jonathan S. Turner, "Design of an integrated services packet network," *Proc. of the Ninth Data Comm. Symposium*, pp. 124-133, Sept. 1985. Also, to appear in the *IEEE Jour. on Selected Areas in Comm.*
[28] Jonathan S. Turner, "Design of a broadcast packet switching network," *Proc. of Infocom '86*, pp. 667-675, April 1986. Also, to appear in the *IEEE Trans. on Comm.*
[29] Jonathan S. Turner, "The routing problem in broadcast packet networks," Washington University, Computer Science Department, technical report in preparation.
[30] Jonathan S. Turner, "The congestion control problem in broadcast packet networks," Washington University, Computer Science Department, technical report in preparation.

Jonathan S. Turner received the B.S.C.S. and B.S.E.E. degrees from Washington University in St. Louis in 1977. He received his M.S. and Ph.D. degrees in Computer Science from Northwestern University in 1979 and 1981.

From 1977 to 1983 he worked for Bell Laboratories in Naperville, IL, first as a Member of Technical Staff and later as a Consultant. His work there included the development of maintenance software and design of system architectures for telephone switching systems. From 1981 to 1983 he was the principal system architect for the Fast Packet Switching project, an applied research project which established the feasibility of integrated voice and data communication using packet switching technology. He has been awarded eleven patents for this work.

He is now an Associate Professor of Computer Science at Washington University, where he continues his research on high performance communications systems. His research interests also include the study of algorithms and computational complexity, with particular interest in the probable performance of heuristic algorithms for NP-complete problems.

Dr. Turner is a member of the IEEE, the Association for Computing Machinery, and the Society for Industrial and Applied Mathematics. ■