# Prospects and design choices for integrated private networks

by P. E. Green, Jr.
D. N. Godard

***This paper reviews some of the choices that will be available in the next few years, as the much-discussed move toward implementing voice and data integration within a single wide-area integrated private network proceeds. After the term* wide-area integrated private network *has been defined, a discussion is given of requirements the network ought to satisfy for its users. Then two particularly promising approaches, fast packet switching (FPS) and hybrid switching (HS), are defined, and specimen design points for FPS and HS are postulated, so that the two can be compared. While a definitive comparison would require systematic cost and performance studies, much insight can be gained from the qualitative comparison that we present here. We assess some of the arguments that have been put forward in favor of FPS or HS and conclude that, although today both architectures have promise, and research on both should continue, FPS appears to be slightly simpler to implement and operate.***

Today's corporate telecommunication networks are composed of at least two stand-alone networks. One network is dedicated to voice—and possibly facsimile and teletype applications—and the other network is dedicated to data communications. Each network has its peculiar interfaces, transport, and switching mechanisms. A widely recognized objective is to combine these switched services into the same network, which should result in lower equipment and communications costs, higher level of service, and greater flexibility for introducing and distributing new services.

The Integrated Services Digital Network (ISDN) concept[1] is the most comprehensive step in this direction. The ISDN standards effort defines the different classes of service that a network may provide. ISDN also defines a unique integrated interface for access to these services. However, it does not dictate how the network transport and switching mechanisms are to be architected and implemented.

In this paper we examine these transport and switching questions for wide-area integrated private networks of the 1990s. This emphasis means that we shall not deal with questions that are specific to local-area networks only, nor with the many problems that are specific to integrated common-carrier facilities. Common-carrier-integrated networks are in many ways a more complicated issue than private networks. The reasons for this are the need to coexist with and migrate from the extensive installed base of circuit-switched voice networking technology and the fact that the volume of traffic is much greater for a carrier than it is for a single business.

Several basic switching approaches may be considered when designing an integrated network to carry voice, data, and some image. These approaches include circuit switching, packet switching, and hybrid combinations of the two. In this part, we focus our attention on the following two particularly promising approaches:

- *Fast packet switching* (FPS),[2] which retains the main characteristics of conventional packet switching but emphasizes the exploitation of high-speed switching and transmission technology and

Figure 1 A typical private wide-area network of the 1990s

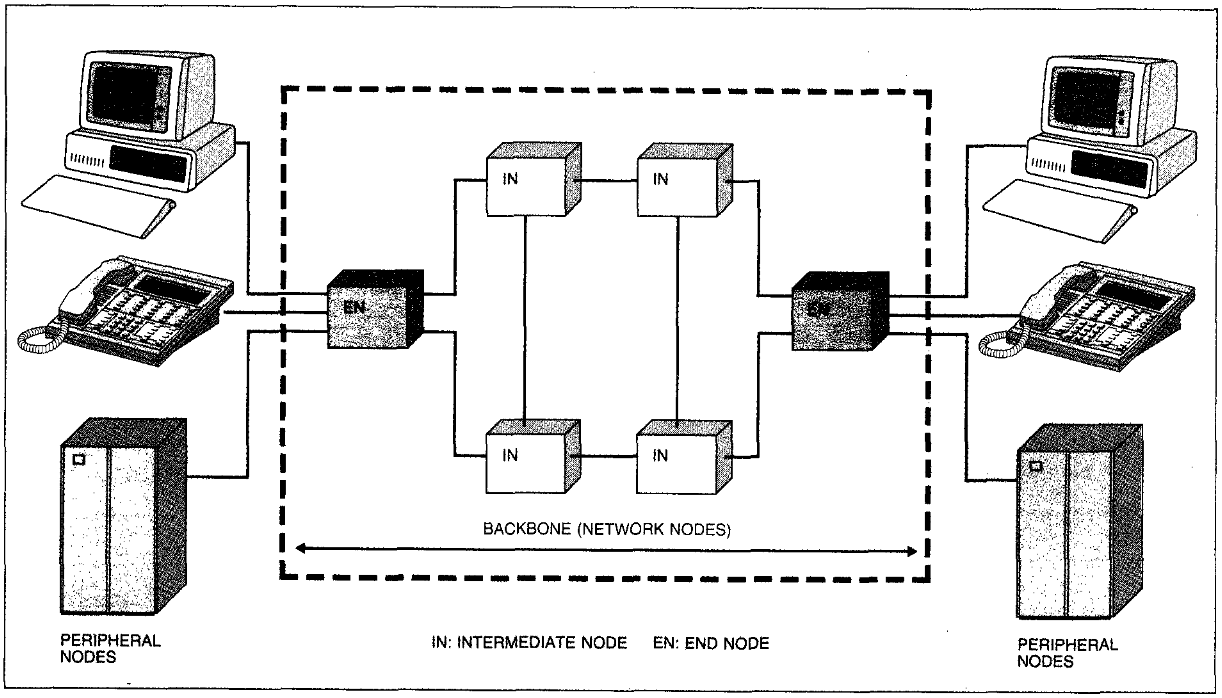


minimizes intermediate-node processing for flow control, routing, and error recovery.

- *Hybrid switching* (HS),[3] in which conventional circuit switching is used for voice and voice-like traffic, whereas fast packet switching is used for data traffic. One way of implementing HS is to have two different switching architectures and different network architectures. Another way is for voice and data to share the same transmission and switching facilities under the guidance of a common network architecture, even though the voice and data would be processed differently in terms of details of buffering, blocking, etc. The second approach is considered in this paper.

Although the focus is on these two approaches, mention is also made at the appropriate point of two other switching methods, known as *fast circuit switching*[4] and *burst switching.*[5]

First, a definition is given of what pieces might make up a corporate integrated telecommunications network. A set of realistic requirements for such a structure is then defined, and sample design points for FPS and HS are postulated, so that the two may be compared. The comparison leads to identification of deficiencies possessed by one or the other. Removal of these deficiencies is a challenge for future invention and research.

## Wide-area integrated private networks

We now consider the model of a future wide-area integrated network for a typical corporation. Such a network, illustrated in Figure 1, is made up essentially of two parts. The first part is a *backbone* consisting of *network nodes* interconnected by high-speed trunks. The second part is the collection of *peripheral nodes* that connect into the backbone by lower-speed *subscriber lines.* The peripheral nodes are the users' installed base of computer terminals, controllers, processors, telephones, and other terminal equipment. There might be from less than a dozen to several hundred or even a thousand network nodes in the backbone. The trunks might number up to ten per network node, and the local attachment of the peripheral nodes, i.e., *lines,* might number up to thousands per network node.

Network nodes are of two kinds. *Intermediate nodes* can do intermediate routing, and they support only trunk interfaces. *End nodes* can do both intermediate routing and line attachment; thus they support both trunk and line interfaces. Peripheral nodes are not part of the backbone and do not participate in the intermediate node routing process.

End nodes might or might not be running applications, in addition to executing network-node function. In another variation, an end node could be split into two physically separated pieces connected by a high-speed link, one piece supporting the lines and the other piece supporting the trunks.

In facing the design choices for the backbone network, some of the elements in Figure 1 must be assumed given, whereas others may be varied. The technology and performance parameters available for the trunks are defined by the common carriers. The subscriber lines are provided by carriers, although some of them may take the form of metropolitan-area networks (such as connections provided by CATV franchisers) or local premises connections, such as local-area networks. Like the trunks, the subscriber lines are not easily changed. Thus, the design choices to be made by the provider of the private network concern the design of the network nodes individually and their organization into a collective network architecture.

New ways of designing integrated networks are now possible, using emerging new technologies. Trunk speeds are going from today's T1 rates (1.5 megabits per second) using twisted pair locally and microwave relay for long haul to rates of hundreds of megabits per second both locally and long-haul by the use of optical fibers. Satellite trunks with multimegabit rates, but undesirable propagation delay characteristics, are also being supplanted by optical fibers, especially in high-traffic backbone situations. Faster switching is being developed, based on silicon today but with the expectation of even higher-speed gallium arsenide circuitry within a few years. The largest switches today already have throughputs in the multiple-gigabit-per-second range, and this maximum is expected to increase. Eventually, fully photonic switching of optical signals from optical fiber lines and trunks will be practical, and thus as signals flow across the network nodes, the signals will never be converted from photons to electrons and back again but will remain in optical form.

The availability of high transmission speeds and nodal throughput speeds has important consequences, not only for the sorts of services that can be provided to the user, but also for the network design itself. For example, end-to-end delay can already be made so short that operations that formerly had to be conducted on a hop-by-hop basis can now be conducted end-to-end. An example of this is data link control error recovery using automatic repeat requests.

**Availability is a prime requirement.**

Several issues that must be faced in designing the backbone are discussed in this paper. The principal issue is the choice of basic form of switching: should packet switching or some hybrid combination of packet and circuit switching be used? A subsidiary question concerns the method of network control; i.e., should it be centralized or decentralized?

### Requirements

In order to decide among the various options available for designing the backbone of a wide-area integrated private network, it is necessary to be clear on exactly what such a network should do.

**Availability.** For integrated private networks of the 1990s, such as that in Figure 1, availability is a prime requirement. The availability issue is one of ensuring the provision of a connection when requested, and of maintaining that connection in the presence of possible failures. For many businesses, availability is being ranked today as often the most important system requirement, even ahead of cost.[6]

The first aspect of availability is whether a requested call can be set up. Call setup is the process by which the logical object at the source peripheral node originates dialing information that the backbone translates into a completed end-to-end path to the desired target logical object in some target node. When this is determined to be impossible, the call attempt is rejected; i.e., the call is *blocked*. Thus, a key availability parameter is the maximum probability of blocking a new call setup during the peak busy hour. We have adopted the one-percent figure widely used in telephone system design[7] as the requirement for both voice and data.

The second aspect has to do with continuous availability of the call, once it is set up. This is expressed by the probability of failure of an existing call. To reduce this number to as low a value as possible, it is useful (for both voice and data) to provide *nondisruptive route switching* if possible. This term refers to the process of rapidly diverting traffic on a failed connection onto another backup connection in such a way that the two end users do not have to experience a new connection setup.

**Integrity of voice and data messages.** Individual messages may be lost accidentally, for example because of transmission errors on the trunk. Other causes might include buffer overflow in nodes where messages are accumulated in buffers of finite size, or whenever messages are lost on purpose, for example as one method of congestion control. Voice users can tolerate occasional short dropouts of the speech signal as long as the probability of such loss stays within prescribed limits. The permissible maxima have been standardized in a number of places. Roughly speaking, the various detailed standards[8] specify that there shall be no more than 0.5 percent probability of dropout ("clipping") of greater than 15 ms duration at both the edges and the middle of talkspurts. Unlike voice, the loss of data messages, once the connection is established, is intolerable, and the standard error-recovery techniques, based on error detection and retransmission, are required. As mentioned earlier, these can be used on an end-to-end rather than on a hop-by-hop basis.

**End-to-end delay.** There is no clear answer to the question of how much delay to allow over and above the propagation delay (25 ms one way over carrier facilities coast-to-coast in the United States[9]). Long delays are an issue in satellite links and also in classical low-speed, packet-switched services based on modest node throughput and trunk speeds. One of the most pernicious effects of delay for voice traffic is the problem of echoes, which arise whenever there are impedance mismatches involving the far-end subscriber line using two-wire technology, a technology that is expected to persist in the installed plant indefinitely. As the round-trip delay of the backbone becomes longer, the subjective annoyance level of a given amplitude of echo signal increases.[10] By making widespread use at end nodes of modern echo-canceler technology,[11] it is possible today to control echoes to a low enough level of audibility, even for the 250-ms one-way delay that would exist if a satellite link were used. The effect on a conversation of the bulk delay then enters the picture in a complex psychophysical way.[12] Although the existing CCITT standards[13] speak of several hundreds of milliseconds as being tolerable, apparently a re-examination of the suitability of satellites for voice is taking place, and it appears quite likely that competitive integrated networks of the 1990s will have to obey considerably stricter delay standards. For present purposes, instead of allowing delays of the order of satellite delays, we shall use a value of 160 ms, i.e., twice the CCITT's number of 80 ms maximum one-way propagation delay[13] for terrestrial voice circuits spanning two countries. It is assumed that none of the trunks are satellite connections.

**The voice quality of the integrated network should be as good or better than toll quality.**

For data, the pertinent question about delay concerns how much more productive an interactive session will become as the delay seen by the user is reduced to small values. A recent study[14] shows that, for many applications, user performance is considerably degraded for response times of several seconds compared to that achievable at less than 500 ms. Therefore, 250 ms is assumed as an upper limit on one-way delay for interactive data.

**Voice quality.** There is some debate about whether today's *toll-quality* voice standards[8] in telephone networks will be adequate for the indefinite future. In the new competitive environment, it seems likely that those price-competitive offerings with higher voice quality will be preferred. For purposes of this paper, the requirement will be adopted that the integrated network should be as good or better than toll quality.

**Complexity of implementation.** In comparing approaches to integrated network design, the intrinsic complexity of the software and hardware implementation at each node must be considered. In addition to this, however, the complexity of understanding and managing the network, once it is built, must also be taken into account.

**A few standard bit rates will become dominant.**

**Trunk utilization.** It is a somewhat controversial question just how important a given percentage difference in throughput will be in designing the network. One school of thought argues that the new technologies (especially fiber) are driving down the fractional system costs represented by transmission to levels that are sufficiently low that maintaining high trunk utilization is no longer very important. On the other hand, all indications are that users will continue to invest in methods of avoiding large transmission inefficiencies. These methods have typically included dial-up service, statistical multiplexing (e.g., with packet switching or multidropping), low-rate voice encoding, interpolation of data into speech gaps, time-assigned speech interpolation (TASI) (which we shall discuss subsequently), and similar measures. It is significant that, according to current ISDN plans, users (including builders of private networks from ISDN trunks, as in Figure 1) will be charged by the bit, independently of distance and bit rate, as they are today for packet carrier service. There will be two consequences of tariffs that are based principally on the number of bits transmitted. For voice, there will be pressure to invest in the silicon necessary to do speech encoding to lower bit rates. For data, the pressure will be for high-speed lines, since users will not want the poorer response time and throughput of a lower-bit-rate service that is no cheaper. A reasonable judgment is that, when comparing two network designs, only if the difference in trunk utilization exceeds ten percent will it be considered significant.

**Arbitrary bit rates of attached lines.** ISDN envisions that a few standard bit rates, such as 16 kilobits per second (D-channel), 64 kilobits per second (B-channel), and so forth will become dominant. Unfortunately, if this ever becomes the case it will only be after many years of living with the present high degree of variegation. For some years it will be necessary to support attachments ranging from today's speeds to the latest in trunks (e.g., optical fibers) and lines (e.g., high-speed local-area networks—LANs). The leading-edge trunk and line technologies will become geographically pervasive only slowly, and meanwhile the lower-bit-rate subscriber connections and trunks are likely to remain in widespread use for a long time. An example of the latter is T1 at 1.544 megabits per second (2.048 megabits per second in Europe). Examples of subscriber line speeds that will need to be supported indefinitely include terminals and heavily compressed voice (around 8 kilobits per second), Adaptive Differential Pulse Code Modulation (ADPCM) voice (32 kilobits per second), PCM voice, ISDN B-channel, and heavily compressed conference video (64 kilobits per second), T1, and conference video (1.5–2 megabits per second), local-area networks (2–10 megabits per second), computer input/output channels (about 30 megabits per second), and T3 multiplexed voice at 44 megabits per second (34 megabits per second in Europe).

Two comments should be made about our list of subscriber line-speed requirements. First, note the absence of broadcast quality video, even though point-to-point and one-to-several conference video are included. Video broadcast service was excluded from our list of requirements on the grounds that the extremely large bit-rate volume and the peculiarly one-to-many topological nature of digital broadcast video means that it is best left to a separate service. This is what is planned in Nippon Telephone and Telegraph's Integrated Network System.[15] Second, something needs to be said about the relative proportion of voice/video and data traffic. We shall assume here that overall voice/video traffic on the backbone will dominate data traffic by a factor of at least five to one in aggregated bit rate and ten to one for many situations.

**Fast connect/disconnect.** The speed with which a virtual circuit is set up for data, or a real circuit for voice, should be as fast as possible. Certainly the setup time should be no slower than that afforded by today's technology. It was mentioned earlier that fast switching and transmission technologies allow reconsideration of familiar assumptions about, for example, per-hop DLC error recovery. The same is true of call setup. A network that allows one-way, end-to-end delays of 160 ms at most may provide new possibilities for rapid call setup and clearing, and these possibilities should be exploited.

**Growth and change.** A large private network should be able to accommodate classes of traffic and magnitudes and bit rates of traffic offered by new applications with the maximum flexibility possible. Telecommunication networks, once installed, are widely known to have a much longer lifetime between upgrades than, for example, the hardware in a typical computer center—even more so with respect to software upgrades. It is relatively easy to upgrade software corresponding to the upper protocol levels as system requirements evolve, but it is considerably more difficult to change the hard-wired and microcoded implementations of the physical layer, which is where much of the switching and transmission expressed in Figure 1 take place.

**Migration and coexistence.** Any new form of network will have to be compatible with the very large installed base of circuit-switching technology represented by the world's existing telephone plant. This is true for both voice and data. In addition, for data, there is an analogous but less severe compatibility problem posed by the installed base of such widely used designs as Systems Network Architecture (SNA), Digital Network Architecture (DNA), Transmission Control Protocol/Internet Protocol (TCP/IP), and X.25, as well as the emerging Open System Interconnect. All these conditions, which will only be more severe by the time of deployment, must be taken into account in deciding among alternatives for tomorrow's integrated private networks.

### Fast packet switching

In order to allow fast packet switching and hybrid switching to be compared with respect to the requirements, we now postulate arbitrary design points for both of them that seem to reflect the current state of the art. In both cases we define what the term means, summarize how network control is to be done, and show how voice and data are to be handled.

**Definition.** The fast packet switch design point assumed is essentially that described in Reference 2, with the addition of some recently developed ideas about decentralized network control.[16] The term *fast packet switching* refers to the exploitation of classical packet switching in a high-speed technology environment. The use of high-speed trunks and low-delay switches allows certain protocol complexities of classical packet switching to be avoided. For example, one may dispense with some of the complex processing of individual packets at all nodes along the path in order to do routing, to control flow and congestion, and to maintain message integrity. At the entry end node, individual packets are assembled from the bit streams arriving from separate attached lines and are buffered awaiting transmission on the high-speed outgoing trunks. At each remaining node along the path, packets are similarly queued awaiting service by the appropriate output trunk. At the final end node, the packets are routed to the appropriate line adaptor for transmission at the assigned bit rate to the peripheral node. A packet length of about 1000 bits and a header length of about 100 bits are assumed, so that there is a built-in loss of bit efficiency of about ten percent.

**The path along which the packets flow is a latched-in or frozen virtual circuit.**

**Control.** The path in the network along which the packets flow is a *latched-in* or *frozen* virtual circuit (VC). That is, successive packets that flow between two logical objects in communication with each other always flow along the same physical path for the lifetime of the VC, and this physical path is the same in both directions.

For various reasons, it is desirable to apply the decentralized form of network control to establish and maintain virtual circuits. Experience has recently been gained in this form of network control with the recently announced Advanced Peer-to-Peer Networking (APPN).[16,17] In this approach, each network node maintains a constantly updated *topology data base* containing the current layout and characteristics (e.g., capacities and delays) of the various trunks making up the backbone. The topology data bases of all nodes are updated whenever a topology change occurs (addition or deletion of trunk or node) by sending *topology-update* messages around the network according to a certain fail-safe and minimal-traffic distributed algorithm.[16]

The topology data base provides each source end node with the information necessary to do the call setup by itself along some desired route. When either a voice or data VC is requested terminating in a logical object associated with some particular target end node, the source end node looks at the topology data base, computes a good route, builds a header containing a field that defines the route as a succession of outgoing links at successive nodes along the path, and sends an exploratory packet to the far end and back. If this route is not acceptable, the next best route is tried. If there is no good route or the target resource is busy, then the call is blocked. The preparatory packet and each succeeding packet on a successfully established VC are internally routed at each node to the outgoing trunk specified in the header.

One may also decentralize the preparatory directory lookup that tells at which physical node a logical resource is located. If the source node does not already know which physical node is the target node, it can initiate a distributed search for it.[18]

Nondisruptive route switching can be carried out as follows. When the route to be used is first computed at the source node, a stack of several backup routes is also computed, and their headers are built and stored. An obvious candidate for a backup route would be one that is topologically disjoint with the primary route. When failure of the primary route is detected at the source node, the next header on the stack is immediately substituted. The receiving end node is expected to follow suit so that the reverse direction of the VC follows the new physical path.

There are felt to be many advantages to the decentralized mode of network control. The avoidance of the single- (or few-) point-of-failure condition is certainly one, but at least as significant is the ability of the network to be reconfigured without significant software and table reloads in the nodes. The possibilities for increasing the speed of call setup and of nondisruptive route switching by having routing information prepared in advance at each originating network node are particularly interesting, and they will be discussed later in this paper. If requests for call setup had to travel to special control centers and back to all nodes along the path before even the exploratory route test packet could be sent to and reflected from the target end node, route setup would take somewhat longer.

Figure 2 Typical network node implementation

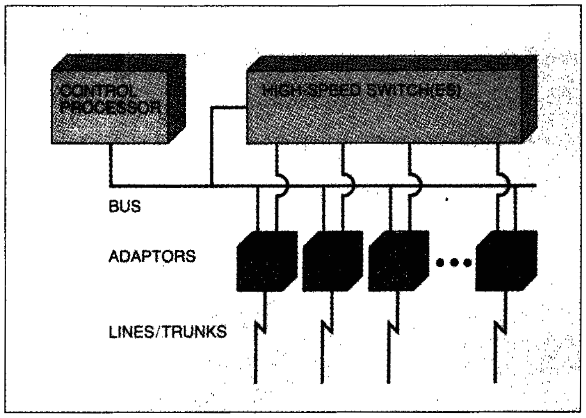


For both the FPS and HS design points, it is assumed that the particulars of subscriber line signaling protocols are beyond the scope of this discussion, because up to a point they are common to both approaches. In both the FPS and HS cases, the node—whether an end node or an intermediate node—consists of one or two high-speed switches (one for FPS and two for HS), a control processor (controller), an assortment of line interface cards (adaptors) that are possibly able to provide some multiplexing and low-speed switching, and some sort of high-speed internal communication structure, such as a system bus, that interconnects all these elements. The block diagram of Figure 2 shows one possible node architecture. Subscriber line signaling at the end node consists of a series of interactions between adaptor and controller, the ultimate result of which is to invoke the establishment of a virtual circuit across the backbone (for FPS) or a real circuit (for the circuit-switched portion of HS). The details of adaptor-to-controller interaction are irrelevant to the comparison we wish to make, but details of the controller-to-switch interaction and the interactions between controllers in the backbone nodes are important and are discussed later in this paper.

Modularity and commonality are two important implementation objectives in both FPS and HS. The network nodes of a network of the 1990s, such as that in Figure 1, should be designed so that they are available in a wide range of total throughput sizes to accommodate a range of user traffic volumes. Also,

**A variety of trunk and line speeds must be supported.**

for a given network node capacity, a variety of trunk and line speeds must be supported. Growing and changing the network will be facilitated in this way. It is also desirable to have a common implementation for intermediate nodes and end nodes as far as possible. The principal difference between the internal architecture of an intermediate node and that of an end node should be that the latter supports not only the small number of high-speed trunk adaptors but in addition a large number of lower-speed line adaptors. The bus-oriented structure of Figure 2 supports the objectives of modularity and commonality for both FPS and HS.

**Voice.** Voice and conference video are served by assembling at the first node packets of about 1000 bits duration and then sending them to the user at the final end node over a virtual circuit, the delay of which has been so controlled that all packets experience the same delay. This requirement for uniform delay is not present in the case of data. Among other differences is the fact that error control and other matters are handled differently for voice VCs and data VCs.

Before voice packets are formed, a speech activity detection operation is performed on the speech signal. Packets are generated only during active speech intervals, a step that is considered necessary because of the large utilization saving obtained by not transmitting silences. A factor of more than two in utilization is available in principle, because most of the time only one of the two conversing parties is speaking. Moreover, during periods of talking there are gaps between sentences, words, and syllables.

Packet switching has the ability to serve attached lines of arbitrary bit rate, but there is one special provision that needs to be made for low-bit-rate voice attachments. In order to minimize the audible effects of losing voice packets and to keep voice packetizing delay to reasonably low values, the buffering of voice bits for packetizing at the first node is limited to about 30 ms. This means that low-rate digital voice subscribers using rates less than 32 kilobits per second will be using packets shorter than 1000 bits. There will then be a greater than ten percent loss of bit efficiency because of the use of a larger fraction of the packet for the header. This design point is felt to be reasonable, because if low-rate voice traffic ever becomes a large component of the total voice traffic, transmission costs are likely to have dropped at least as much by that time.

At each node along the route, packets must be buffered awaiting transmission on an outgoing trunk. To minimize voice end-to-end delay, voice packets are given a higher priority than data packets. The buffer management scheme must be designed so that the maximum permitted queue length before overflow occurs is made shorter for voice packets than for data packets. This is the case because, with voice, minimizing path delay is more important than preventing the loss of an occasional packet, whereas the converse is true for data.

Special corrective means must be provided to ensure that the delay that the users see after the packet has passed across the backbone is constant and as small as possible,[19] in spite of the variability of buffering delays of packets in the nodes along the route. One procedure is illustrated in Figure 3. At the entry node a special "accumulated time delay" field in the header of each packet is initialized to zero, and as the packet passes through each successive node, this field is incremented by the observed intranode delay. Each node measures its own intranode delay using only a crude local clock. In the figure, the total accumulated nodal delay reading at the exit node for the $i$th packet is shown as $D(i)$. Thus, by using this scheme, the delays of the various packets arriving at the final node will have been measured to within a constant equal to the unknown sum of all the trunk transmission delays. This constant, $C$ in Figure 3, which will not vary during a call, can be computed roughly from knowledge of the topology. At the exit node, a variable amount of delay $V(i)$ is deliberately added to the observed accumulated delay $D(i) + C$ so as to equal a total delay $T$. Occasionally, it will happen that $D(i) + C$ is already greater than $T$. That is, the actual time of flight of the $i$th packet exceeds $T$. In such cases, the packet is discarded. $T$ is periodically readjusted so as to maintain the end-to-end delay as small as possible, while never allowing the frequency of packets lost by deliberate discard plus buffer overflow to exceed one percent. The maximum value that $T$ is allowed to attain before starting to block voice calls is 160 ms.

**Data.** Although congestion control for voice packets is a matter of discarding overflow packets for existing calls and of blocking calls when the overflow rate becomes excessive, a different procedure is appropriate for data. Deliberate discard of packets cannot be tolerated. Flow control may be exerted between the end nodes on each data virtual circuit or on individual hops on the VC. It may be possible to use only end-to-end flow control if propagation delays are small, because the high transmission speed of the system allows it to react quickly to local congestion along the route and to do so from the route ends. The particular way in which flow control is exercised, on either an end-to-end or a per-hop basis, is to use a variable-window (number-of-permits) scheme such that the receiving end tells the sending end how many packets it is now allowed to send in addition to those already sent.

The integrity of data VCs must be maintained by the usual method of error detection followed by requests for retransmission. Packet loss may occur because of line errors or buffer overflow. As with flow control, this error-recovery procedure can be carried out on an end-to-end basis rather than hop-by-hop and still maintain rapid response to congestion conditions, again not only because of the high speed of the system but also because of the low error rates expected with high-speed trunks, compared with the error rates that are common today with subscriber lines. Data packets, which are all of lower priority than voice, themselves are assigned one of several internal priority classes, interactive traffic being given a higher priority than batch traffic, for example.

## Hybrid switching

**Definition.** The second approach for integrating voice and data into a single communication network is *hybrid switching*,[3] in which circuit switching is used for voice connections and packet switching for data transmission. Integration may be achieved in different degrees, the ultimate one being reached when voice and data not only access the transport network through a single interface, but also share with network control information the same physical trunks, and are processed by the same multiplexing and switching equipments under the guidance of the same network management architecture.

Figure 3 Time reassembly of successive voice packets to form a constant-delay virtual circuit

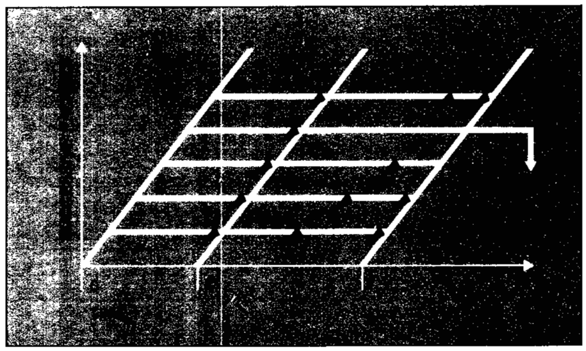

The fact that circuit switching[20] is used for voice implies that synchronization has to be maintained across network nodes. It is assumed that network nodes will derive their timings from the trunk signals, for which synchronization is provided by the common carrier. The conventional concepts of framing and time division multiplexing apply. That is, the bit stream on each trunk arrives in repetitive *frames*, and different *slots* within each frame may be allocated to different connections. We shall assume that the slot widths are multiples of one bit, so as to support subscriber lines at integral multiples of a fairly low bit rate. Specifically, for the standard 125-microsecond frame length of T-carrier and ISDN designs, trunks whose speeds are integral multiples of 8 kilobits per second can be supported without intermediate levels of multiplexing.

The hybrid switching design point considered in this paper assumes that circuit and packet connections are sharing trunk capacity in two ways, as shown in Figure 4.

First, for each frame on the trunk comprising, say, $N$ slots, there are a given number of slots, say $M$, reserved for packets, i.e., slots that voice circuits cannot use. Packets appear sequentially, not interleaved, so that one packet might be spread over the packet capacity of several successive frames. Second, circuit-switched connections have priority over the $N - M$ remaining slots, but packet traffic is allowed to "grab" those of the slots that are left unused at any moment by circuit-switched traffic. In the figure a portion of the circuit capacity is shown being

Figure 4 Frame structure for movable boundary hybrid switching

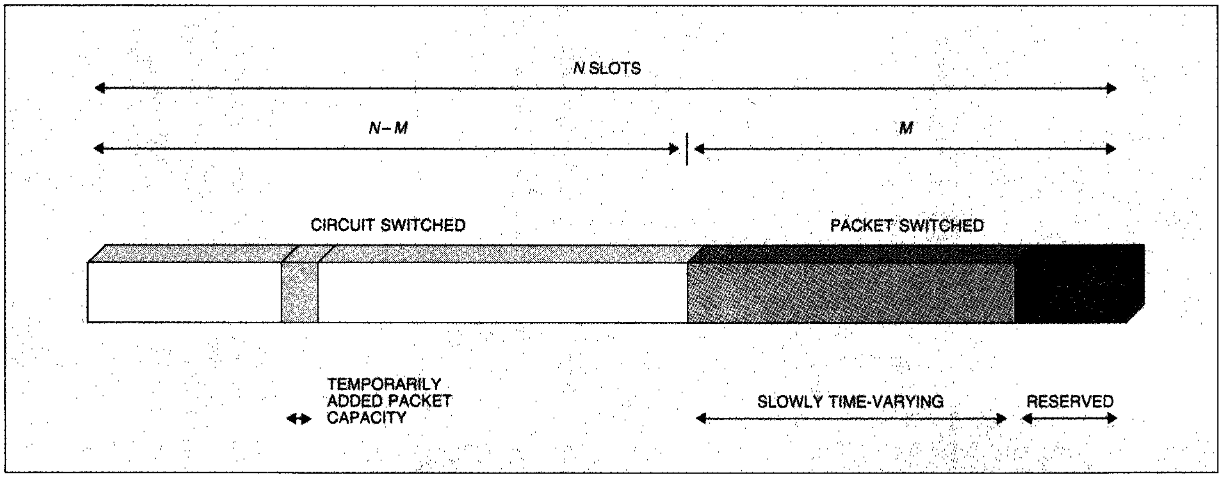


grabbed as "temporarily added packet capacity." Because the relative proportion of voice and data traffic may vary from one node to another, or as a function of the hour of the day, multiplexing and switching equipment must be engineered in such a way that $M$ will be a changeable system parameter.

Many systems being implemented today as part of ISDN may be viewed as HS systems with the boundary fixed. (See, for example, Reference 21.) They embody separate packet and circuit switch fabrics and share trunk slots between packet and circuit traffic in a fixed pattern.[22] It is possible that future HS systems will continue to use fixed boundary schemes, because, although there are utilization advantages to making the boundary movable and allowing grabbing, there is a price in the form of added control complexity. However, in the present study we assume $M$ to be a variable and grabbing to be permitted, because we are trying to compare the best fast-packet-switch case with the best hybrid-switch case.

**Control.** The same decentralized control principle is assumed for the HS design point as was used for FPS. For example, route determination is done by the source end node using its topology data base. The control information that has to be exchanged between network nodes for updating the topology data base as well as for call setup is transmitted by packets. Thus, control packets are handled in exactly the same way as in the FPS case. The procedure for establishing a voice circuit is also very similar to that used for setting up a virtual circuit. The source end node sends an exploratory packet, the header of which contains routing information, as well as the capacity requested for the call. As this packet flows through the network, it reserves the capacity required so that it cannot be used by packet traffic or by another circuit. Control packets are echoed back to the source end node, indicating that the connection has been established, or that the destination is busy, or that the requested capacity is not available on a given trunk. In the latter cases, the packet that is echoed back releases the trunk capacities that had been reserved so far, and another route is tried. If no routes are available or the destination has been found busy, the call is blocked.

As with the FPS design point, control traffic in the HS design point is assumed to flow in packets over virtual circuits between the controllers in the various nodes. A bare minimum packet capacity, shown in Figure 4 as $M_{\min}$, must be reserved for that purpose.

The partitioning of capacity between packet and voice traffic, as shown in the figure, is managed on individual trunks by the interaction of the controllers at the two trunk ends.

**Voice.** A constraint of the circuit-switching approach is that the bit rate presented by the end user must be a submultiple of the trunk bit rate. If, for the sake of bandwidth conservation, low-bit-rate sources are to be attached to the network (e.g., those using vocoders at or below 1–2 kilobits per second), additional adapters carrying out bit-rate matching and multiplexing are necessary in order to build up the bit rate to fill out at least one bit per frame.

Another characteristic of conventional circuit switching is that the bandwidth associated with a call is made available continuously during the whole duration of the call, whether active speech is present or not. Instead, we shall consider here a more elaborate definition of circuit switching designed to exploit the gaps in human speech, so as to optimize trunk bandwidth utilization, as in the case of FPS.

In hybrid switching, it is possible to take advantage of the pauses in speech either to transmit data packets during the silences (using grabbing, already referred to) or to interpolate additional talkers onto the trunk. The latter technique, called digital speech interpolation (DSI),[23] is the digital equivalent of analog time-assigned speech interpolation (TASI).[10] The way TASI and DSI work is to fill time intervals in the outgoing trunks that would otherwise contain only silence with the "talkspurts" from various incoming lines. Enough incoming lines must be aggregated so that the probability of finding no empty place to put a new talkspurt is low. In this way, one may often fit a number of (intermittent) incoming speech circuits onto a number of (continuously occupied) outgoing circuits equal to as little as half the number of incoming circuits.

Given our projection that overall voice traffic on the backbone dominates data traffic, there is a strong incentive to favor efficiency in handling voice traffic at the expense of data. Therefore, the use of DSI in hybrid switching will be assumed in comparing HS with FPS. The grabbing of circuit capacity for packet traffic will occur only if some circuit capacity is left over and not needed for DSI.

A fundamental difference that remains between this extended form of circuit switching (i.e., using DSI) and packet switching is that talkspurts for which no trunk bandwidth is available[24] at the time they are produced are frozen out instead of being queued.[25] A second difference is that the DSI functions are performed at the trunk level, and not only at the end-user subscriber line-adapter level.

**Data.** Data are handled by the HS network in exactly the same way as by the FPS network, as described earlier in this paper. It is worth mentioning in addition, that congestion control must take into account the fact that the overall bandwidth allocated to packet traffic is made variable.

## Specific comparisons

In this section we compare the capabilities of the two forms of switching to satisfy the various requirements discussed earlier in this paper.

**Availability and integrity.** Fast packet switching has an advantage over hybrid switching in the relative ease with which nondisruptive route switching may be implemented for voice circuits. The circuit-switched component of HS uses dedicated slots to carry real circuits. If nondisruptive route switching were to be implemented for real circuits, the backup capacity would have to be preallocated. That is, the computation to identify the slots would have to be made beforehand. Without hop-by-hop DSI, the specific slots are held in reserve. With DSI, the required number of slots are held in reserve. Either way, this is an unsatisfactorily wasteful procedure. With FPS, a stack of routes and the associated routing headers may be prepared in advance at the source node, and the next one popped from the stack when needed. For data connections, FPS and HS have comparable availability, because packet-switched VCs are used in both cases.

**With digital speech interpolation, the required number of slots are held in reserve.**

As for message integrity, hybrid switching would have an integrity advantage for voice. This would be true if DSI freezeouts were absent, in that none of the voice waveform would be deliberately thrown away, as is done with occasional FPS voice packets, either from buffer overflow or delay overflow. (See Figure 3.) However, DSI is considered important enough for trunk utilization reasons that it was included in the HS design point. Since DSI freezeouts with HS lose about the same fraction of the voice waveform as voice packet discard does with FPS, the two approaches are comparable in this respect.[26]

A minor additional integrity drawback of HS is the occasional loss of one 8-bit TDM sample for all voice and data channels within the trunk service provided by the carrier that will occur due to "slips" of one integral frame length in clocking. This happens only

**Hybrid switching is a better solution from the delay standpoint.**

once per day or less with the clock technology that the common-carrier plant uses today.[10]

**End-to-end delay.** Hybrid switching is certainly a better solution from the delay standpoint, because it is voice traffic that is most sensitive to delay, and real circuits are used for voice—not virtual circuits, as with FPS. Although the 160-ms figure adopted for the FPS design point meets the requirements stated for voice, under heavy load conditions, both this figure and the 250-ms figure for data will occasionally be exceeded.

It should be pointed out that if digital speech interpolation is not used, movable-boundary hybrid switching has one pernicious property that must be reckoned with in designing the flow-control algorithms for data. When the extra capacity for packets that is temporarily available on the circuit-switched side of the boundary (see Figure 4) is reclaimed for circuits, it is reclaimed for a long time, i.e., of the order of minutes. Simulation studies[27] have shown that if this extra capacity had been pooled with the other packet capacity and treated as though both had the same statistical properties as the latter, then, when an attempt is made to reclaim the extra capacity, extremely long packet queues may build up waiting for some circuit capacity to be relinquished. This problem is mitigated when DSI is used, because it is much more likely that grabbable circuit capacity will appear after a short interval (seconds) than if DSI is not used, and the occurrence of long data-packet queues is thus not likely.

**Voice quality.** Overall time delay, which is significantly greater for FPS than for HS, may be considered an aspect of voice quality. With the overall delay figure of 160 ms assumed as a requirement, the perceptual experiments reported in the literature indicate that no noticeable degradation of speech quality or ability to converse will occur, so long as echoes are absent. However, the listener tolerance to an echo of a given amplitude degrades rapidly with echo delay, and at 160 ms, echo control using echo cancelers is definitely required. The data on listener tolerance[10] show that the amplitude of the echo after processing in the canceler can be about 6.0 dB poorer at 160 ms than at 250 ms, the corresponding satellite delay, for the same subjective speech quality.

When low-bit-rate voice encoding is used with FPS, the loss of individual packets is likely to be more audible than with higher bit rates. Even under the assumption in the section on fast packet switching that the packet duration does not increase with lower encoding rate, the effect may still be present for some encoding algorithms. This is true because, as the amount of bandwidth compression in the encoder increases, the decoding algorithm may depend on the availability of a longer and longer preceding sample of reliable speech.

The effect on voice quality of packet dropping with FPS and DSI freezeout with HS was discussed earlier under the topic of availability and integrity.

**Complexity of implementation.** Relative complexity of FPS and HS can only be precisely quantified by designing and implementing both approaches.[28] However, it is worth identifying the fundamental differences between FPS and HS that may have a direct influence on system cost or design complexity.

Two kinds of statement can be made. The first has to do with specific functions, algorithms, and procedures that are required by one approach and not by the other, and that are mostly related to hardware. Examples of these are delay equalization (FPS), global network synchronization, and variable-boundary adaptation (HS). The second kind of statement deals with the differences in network management and network control that are reflected in software complexity.

For maintaining voice quality at a sufficiently high level, FPS requires time-stamp incrementing at intermediate nodes and delay equalization at destination end nodes, as shown in Figure 3. As stated in the requirements section, this delay—mainly due to packetization and queueing delays rather than propagation delay—may be as large as 160 ms. This is not in itself an impediment to good speech quality, on the condition that echoes caused by impedance mismatch at two- to four-wire conversion points are adequately canceled. In an FPS network it is very likely that all two-wire local subscriber loops will have to be equipped with echo cancelers, whereas in a HS network—where the only delay encountered is essentially the fixed propagation delay—such echo cancelers are required only if there are very long routes having a propagation delay greater than, say, 30 ms. Although delay equalization and echo cancellation[11] algorithms may be fairly complex to design, the availability of signal processor chips makes their implementation inexpensive, especially considering that these signal processors may be shared for other functions, such as speech activity detection, speech compression, and the detection of dialing pulses or tones.

**Multiplexing and switching equipments must operate synchronously.**

A more significant complexity element is the global network synchronization required by circuit switching in the HS case. Time-division multiplexing of several digital signals requires that these signals be synchronized and framed. All multiplexing and switching equipments in the HS network must operate synchronously, according to a common reference frequency on a network-wide basis. The problem of providing a very stable reference frequency is left to the common carrier, so that this frequency is derived from an attached trunk. Once all nodes have derived the same frequency, the phase differences on different trunks must be compensated.[29] Even if the problem of providing a very stable reference frequency source for a node is left to the carrier in this way, this frequency has to be recovered from one particular trunk, defined as the master trunk, and distributed down the node's multiplexing hierarchy. Furthermore, means must be provided to switch the derived synchronization as nondisruptively as possible to another trunk in case of failure of the master trunk.

This is in contrast with packet switching, in which the trunk transmission speed is independent of user data rates. All multiplexing and switching operations may be achieved in an asynchronous manner, in intermediate as well as peripheral nodes.

Packet and circuit switching are very different methodologies, and it is difficult to be precise as to which approach is cheaper to implement or easier to design. Just because the methodologies are different, HS can be shown (under conditions of avoiding duplication of buffers) to require two switch fabrics at end nodes and intermediate nodes. On the other hand, FPS requires only a single switch fabric with different queues for handling voice and for data of different priorities. Moreover, since trunk bandwidth is shared by voice and data in a dynamic manner, each of the two switch fabrics in the HS case should be designed so as to be able to handle all the bandwidth that will ever be required of them. Then it may be expected that a significant part of the overall HS switch capacity will be wasted on the average.

The same argument holds true for network management and network control. In the HS network, path setup processes are different for voice and data. Both the "movable-boundary" concept and digital-speech interpolation require that appropriate protocols be implemented between the node pairs at both ends of each trunk for specifying how trunk slots are allocated at any given instant. These DSI and movable-boundary algorithms and protocols are not needed in the FPS case.

The next section gives more detail on why the way that DSI uses pauses in the conversation for conserving trunk capacity adds to the complexity of HS.

One concludes that FPS is somewhat cheaper to build and less complicated to operate than HS. This is because FPS does not need any network-wide synchronization, speech gaps may be more easily exploited, and it uses only one switch fabric and almost the same protocols and procedures for both voice and data. FPS presupposes essentially one network, which should lead to simpler software development.

**Trunk capacity utilization.** The HS and FPS design points given previously assume that digital speech interpolation (DSI) and speech activity detection

(SAD), respectively, are used for optimizing trunk utilization.[30] Theoretically, HS can intrinsically make better use of trunk capacity for the following reasons: (1) FPS traffic must remain below full bandwidth utilization (say at 60–70 percent) to keep queueing delays acceptably low, and (2) HS does not experience the overhead inherent in transmitting the packet headers, although there is a small overhead of a fraction of a percent[23] for the slot-management protocol. DSI exploits silences in voice conversations to approximately double the number of voice calls that can be placed on a given trunk.

In order that there be a reasonably high probability of finding an empty time slot to accommodate a given talkspurt, DSI is applied to a large number of voice inputs simultaneously at the entry node. For example,[31] 70 incoming voice circuits (having speech gaps) can be compressed in this way onto 40 outgoing voice circuits (without gaps) with only 0.5 percent probability of freezeout loss due to unavailability of a slot.

When DSI is applied to a large number of voice calls at an entry end node, in general not all these calls are intended for the same destination end node. This fact makes it impossible, at intermediate nodes, to switch these voice circuits together as they stand. Before entering a switch at an intermediate node, the trunks have to be demultiplexed, synchronous bit streams reconstructed for all of the voice circuits, and, finally, remultiplexed to conform to the switch speed (unless, of course, all the voice circuits on a trunk happen to have the same destination). The process of filling gaps with talkspurts must take place all over again on each output trunk beyond the switch if the DSI statistical advantage is to be gained on that trunk. This is in contrast to the way SAD can be exploited by FPS. SAD, together with nongeneration of packets during silences, is performed once at the entry end node and need not be performed again at each node on the path.

If, therefore, HS is to achieve or surpass the trunk capacity utilization of FPS, it requires DSI equipment at both ends of each trunk at each intermediate node. This difference constitutes a cost advantage of the FPS approach when trunk utilization matters, as it will in a practical network.

**Arbitrary attached bit rates.** As was pointed out in discussing requirements for future integrated private networks, the attached subscriber lines exhibit today a wide variety in speed and in other specifications. One may expect this variegation, which exists for both voice and data attachments, to become worse if anything before there is some eventual convergence to a small number of standards.

One of the advantages of a network based on packet switching is its ability to accept inputs at its periphery of any data rate below that which will saturate its resources. The packets are formed in input buffers at the line bit rate and read out at the trunk bit rate. Because both the FPS and HS design points use packet switching for data, they are on an equal footing in this respect.

**The ability of fast packet switching to accommodate arbitrary voice bit rates constitutes an advantage.**

For voice, the story is different. The ability of FPS to accommodate arbitrary voice bit rates constitutes an advantage of FPS over HS. There is considerable evidence that 64 kilobits per second is not the only voice bit rate that will be popular in the future. It is not just that voice encoding rates of 32, 16, and even 8 kilobits per second can be achieved; they can be achieved with hardware that is becoming very attractive economically and with acceptable speech quality. It is thus important to provide the one-bit granularity assumed for the circuit-switched slots in HS. Even with this feature, HS still does not have the intrinsic flexibility of attachment bit rate that FPS has.[32]

**Fast connect/disconnect.** Decentralized control should have an advantage over centralized control in the faster speed with which a real circuit or a virtual circuit can be set up, for the reasons depicted in Figure 5. The figure shows a three-node backbone network and indicates by numbers the sequence of messages that flow in establishing a real circuit between the leftmost and rightmost nodes. With centralized control, each segment along the route must be set up by an exchange of messages between the segment end points and the central control. Only after this has been completed can the source node be notified that the call has been set up. With decentralized control, as mentioned earlier, the origin node has enough knowledge to properly vector the pathfinder message to the correct target node, where it is sent back again to the source. The individual segments may be set up sequentially. We see that the number of steps is smaller for the decentralized case.[33]

In the design points of both FPS and HS, decentralized control has been assumed for just this reason of fast call setup, as well as to support nondisruptive route switching. However, even with similar control architectures, the time taken to set up a circuit-switched real circuit is likely to be more than the time to set up a packet-switched virtual circuit, for the same reason that nondisruptive route switching takes longer. A packet switch call setup requires only that a viable path be identified, not that capacity be reserved all along the route. Once the entry node receives word that there is a viable virtual circuit to the destination and back and that the receiving peripheral node is not busy, transmission can begin. Later, during use of the path, the routing header will act to define the path as the messages flow; nothing needs to be reserved in advance at the intermediate nodes.

This is the appropriate point to mention *fast circuit switching*, which is often proposed as an alternative to packet switching. One well-known protocol for doing this is the *short-hold* mode of X.21 circuit switching.[4] From the foregoing discussion, it can be seen that such protocols, requiring as they do the connection and disconnection of a new physical path for every message, simply take too long for many practical situations. To repeat, with fast circuit switching, it is necessary not only to check that the path is viable but to also commit the resources.

**Growth and change.** Once the backbone network is installed and begins to evolve further, the need will arise to change the number and speed of lines and trunks attached to each network node and to expand or reduce the capacity of the node in an effort to optimize the cost performance. The relative capability of FPS and HS to accommodate a variety of attached subscriber lines has already been discussed, and it was pointed out that FPS imposes little constraint on the speed of attached lines. The same superior ability of FPS to accommodate a variety of

Figure 5 Comparing the number of messages exchanged in setting up a circuit using decentralized control (top) and centralized control (bottom)

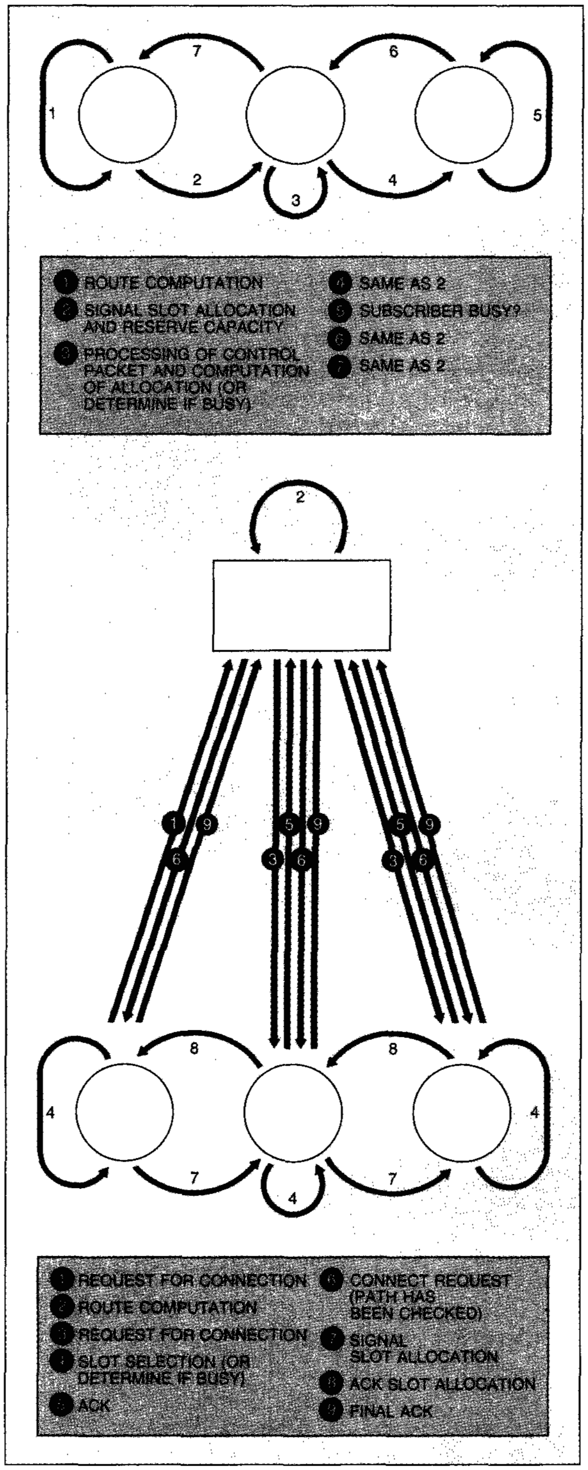

speeds or to change those speeds that are in use applies to the trunks as well as the lines.

The basic reason for this superior flexibility of FPS is that with the circuit-switched portion of HS there is a close coupling of the trunk architecture (frame length, bit rate, subframing, etc.) with the architecture of the circuit switch (number of ports, slot duration, length of time-slot interchange elements, etc.). This is not true with packet switching. Thus, if a number of new trunks are to be added, new adapter cards will have to be added for both HS and FPS. However, with HS the switch configuration will probably have to change (or additional levels of TDM multiplexing will have to be provided), even though the switch might be only lightly loaded. With FPS, the switch need not change as long as it is already fast enough.

If capacity must be increased, with the packet switch a remapping to a faster technology could be a sufficient option; with HS, the switch must be redesigned. As a specific example, if trunks are upgraded (e.g., from T3 to T4), fundamental redesign of the circuit switch is required.

**Migration and coexistence.** One criticism that is sometimes leveled at FPS is that circuit-switch technology is so widely installed that it is unnatural to build a network based on any other kind of architecture. The common-carrier plant carries mostly voice, and it is all circuit-switched. On the subscriber's premises, voice traffic circuit-switched by PBXs dominates data traffic by a large factor and probably always will.

Our conclusions are that this is a minor issue. If FPS is implemented within the backbone of a future integrated network, the fact that it sends packets will be invisible to the external line appearances and insensitive to the way in which the circuit-switched common-carrier lines are used to form the physical-level resources of the backbone trunks. A time-compensated virtual circuit ought to be indistinguishable from a real circuit of the same end-to-end delay.

There is the issue of migration/coexistence not only with existing circuit-switched voice service, but also with existing packet-switched data service, for example, using SNA. Because both the HS and FPS design points use packet switching for data, to first order both would present about the same migration problems with respect to SNA.

## The future: Possible improvements

The comparative examination of fast packet switching and movable-boundary hybrid switching given in the last section has had to be qualitative, in the absence of those quantitative insights that can only be gained from extensive performance modeling, prototyping, and testing. Nevertheless, the discussion has been sufficient to point up several deficiencies of each approach so as to motivate work on further improvement.

It was assumed in the design points that there is a sufficient premium on trunk utilization that with FPS, no packets are generated during speech gaps, and that with HS, digital speech interpolation is used to statistically fill speech gaps in one voice input with talkspurts from other speech inputs. This assumption added considerably to the system complexity of HS.

On balance, it appears that the fast packet switching design point more nearly satisfies the requirements for future integrated private networks than does hybrid switching. The reasons for making this assertion are that with fast packet switching, node implementation and control complexity seem to be less; availability enhancement by nondisruptive route switching is more easily done; call setup time can be made shorter; and there is greater freedom to attach various line speeds. On the other hand, hybrid switching possesses the advantages over FPS of shorter propagation time (leading to reduced echo-control requirement in order to maintain voice quality) and better trunk utilization for voice traffic.

There is clearly room for much research and innovation in improving on the current FPS and HS state of the art, as represented by the design points discussed in the sections on fast packet switching and hybrid switching. One high-priority item is just the question of improving HS trunk utilization for voice without incurring unnecessary complexities with digital speech interpolation on each trunk along a multihop path. Ways of speeding up nondisruptive route switching and call setup for HS also clearly need investigation. For FPS, effective delay-minimizing buffer-management schemes for voice packets and flow-control protocols for data are needed. Ways of minimizing the impact on voice quality of dropped packets or DSI freezeouts are needed. The proper design of adaptive boundary-readjustment algorithms for HS has yet to be understood. Given that HS requires two high-capacity switch fabrics per node, whereas FPS requires but one, economical switch implementations are needed for HS, in order that the switch fabrics will represent a small fraction of the node cost. They should use common part numbers wherever possible and be controlled by unified software.

This paper has focused more on hardware than software. If recent experience within both the computer and the telecommunications communities is any guide, software costs will dominate in the long run. A complete comparative assessment will not be possible until design of the software is dealt with.

## Acknowledgments

We would like to acknowledge many insights gained about fast packet switching from our colleagues Alan Baratz, Inder Gopal, and Bharath Kadaba, and about hybrid switching from Hamid Ahmadi, Israel Cidon, Gautham Kar, and Luke Lien. Jeffrey Jaffe provided many helpful comments about various details of the study. Professor Maurizio Decina's ideas on ISDN and Jonathan Turner's insights into fast packet switching were especially illuminating.

## Cited references and notes

1. Special issue on Integrated Services Digital Network: Recommendations and field trials, M. Decina, W. S. Gifford, R. Potter, and A. A. Robrock, Editors, *IEEE Journal of Selected Areas in Communications* **SAC-4,** No. 3 (May 1986); also companion issue on ISDN equipment development, M. Decina, W. S. Gifford, R. Potter, and A. A. Robrock, Editors, *IEEE Journal of Selected Areas in Communications* **SAC-4,** No. 8 (November 1986).
2. J. S. Turner, "Design of an integrated services packet network," *Proceedings of the Ninth Datacomm Symposium, ACM SIGCOMM Computer Communication Review* **15,** No. 4, 124–133 (September 1985); J. J. Kulzer and W. Montgomery, "Statistical switching architectures for future services," *Proceedings of the ISS,* Florence (May 1984); J. S. Turner and L. F. Wyatt, "A packet network architecture for integrated services," *Proceedings, Globecom* (December 1983), pp. 2.1.1–2.1.6; W. L. Hoberecht, "Layered network protocols for packet and voice data integration," *IEEE Journal of Selected Areas in Communications* **SAC-1,** No. 6, 1006–1013 (December 1983).
3. K. Kummerle, "Multiplexer performance for integrated line- and packet-switched traffic," *Proceedings of the ICCC,* Stockholm (1974), pp. 508–515; P. Zafiropoulo, "Flexible multiplexing for networks supporting line-switched and packet-switched data traffic," *Proceedings of the ICCC,* Stockholm (1974), pp. 517–523; E. Port et al. "A network architecture for the integration of circuit and packet switching," *Proceedings of the ICCC,* Toronto (August 1976), pp. 505–514; H. Rudin, "Studies in the integration of circuit and packet switching," *Proceedings of the ICCC,* Toronto (June 1978), pp. 20.2.1–20.2.7.
4. G. Hellman, K. F. Larsen, and L. Lilja, "Using X.21 switched service: Network, protocols, implementations, and benefits of Datex," *Journal of Telecommunication Networks* **3,** No. 3, 251–267 (1984).
5. S. R. Amstutz, "Burst switching—An introduction," *IEEE Communications Magazine* **21,** No. 8, 36–42 (November 1983).
6. M. A. Fischetti, "Bypassing is big business," *IEEE Spectrum* **22,** No. 11, 78–83 (November 1985).
7. Bell Laboratories Technical Staff, *Engineering and Operations in the Bell System,* Second Edition, AT&T Bell Laboratories, Inc., Murray Hill, NJ (1984); AT&T Customer Information Select Code 500–478.
8. J. G. Gruber and N. Le, "Performance requirements for integrated voice/data networks," *IEEE Journal of Selected Areas in Communications* **SAC-1,** No. 6, 981–1005 (December 1983).
9. M. B. Carey, H. T. Chek, A. Descioux, J. F. Ingle, and K. I. Park, "1982–83 End office connection study," *AT&T Bell Laboratories Technical Journal* **63,** No. 9, 2059–2090 (November 1984).
10. Bell Telephone Laboratories Technical Staff, *Transmission Systems for Communications,* Bell Telephone Laboratories, Inc., Murray Hill, NJ (1982).
11. D. G. Messerschmitt, "Echo cancellation in speech and data transmission," *IEEE Journal of Selected Areas in Communications* **SAC-2,** No. 2, 283–297 (March 1984).
12. P. T. Brady, "Effects of transmission delay on conversational behavior on echo-free telephone circuits," *Bell System Technical Journal* **50,** No. 1, 115–134 (January 1971).
13. Recommendation G.114, "Mean one-way propagation time," *CCITT Red Book* **III,** Fascicle III.1, VIII Plenary Assembly (1984).
14. A. J. Thadhani, "Interactive user productivity," *IBM Systems Journal* **20,** No. 4, 407–423 (1981).
15. T. Murakami, "Inception of INS experience—model system sets in service," *Japan Telecommunication Review* **27,** No. 1, 2–16 (January 1985).
16. A. E. Baratz, J. P. Gray, P. E. Green, Jr., J. M. Jaffe, and D. P. Pozefsky, "SNA networks of small systems," *IEEE Journal of Selected Areas in Communications* **SAC-3,** No. 3, 416–425 (May 1985).
17. R. J. Sundstrom, J. B. Staton III, G. D. Schultz, M. L. Hess, G. A. Deaton, Jr., L. J. Cole, and R. M. Amy, "SNA: Current requirements and direction," *IBM Systems Journal* **26,** No. 1, 13–36 (1987, this issue).
18. I. S. Gopal, A. E. Baratz, and P. Kermani, "The ADVISE directory architecture," in preparation.
19. W. A. Montgomery, "Techniques for packet voice synchronization," *IEEE Journal of Selected Areas in Communications* **SAC-1,** No. 6, 1022–1028 (December 1983).
20. J. C. Bellamy, *Digital Telephony,* John Wiley & Sons, Inc., New York (1982).
21. R. de Hoog and C. Wiebe. "A switching architecture for ISDN," *Proceedings International Communication Conference,* Toronto (June 1986), pp. 24.4.1–24.4.4.
22. Note that the ISDN "$nB + D$" convention, with $n$ a system parameter, is not, strictly speaking, fixed-boundary hybrid switching unless it deals with the trunks, not the user interface, and dedicates the $B$ channels to circuit traffic.
23. S. J. Campanella, "Digital speech interpolation," *COMSAT Technical Review* **6,** No. 1, 127–158 (Spring 1976).
24. For some forms of speech encoding, it is possible to respond to congestion by "bit dropping," deleting the least significant bits. This option is available for both HS and FPS. See References 23 and 34.

25. Queueing is possible, and it will produce further gains in utilization (see Reference 31)—but at the expense of the same time-stamping and reassembly problems that are required with packet switching. Burst switching (Reference 5) is a form of fast packet switching that uses limited queueing of voice in this way and works as follows. Each new talkspurt in an input 64-kilobit-per-second voice signal is designated as a voice packet of talkspurt length. At each node along the path, an attempt is made to seize for that packet an 8-bit slot in recurrent 125-microsecond frames. To avoid large and variable delays, the voice packet queueing that is performed at each node is not allowed to exceed 2 ms. If output trunk congestion causes the delay to exceed 2 ms, then 2 ms of samples are discarded. Time-stamping and reassembly are not required, and only one switch fabric is used.
26. In a recent study comparing burst switching and FPS, O'Reilly (Reference 35) argues that, since the freezeout of DSI and the clipping of burst switching both occur at the sensitive leading edge of a talkspurt, the effect of a given percent loss of this form is more noticeable aurally than the same percent of lost packets, since the latter will occur at random places in a talkspurt and moreover can be partially repaired by splicing and interpolation techniques.
27. C. J. Weinstein, M. L. Malpass, and M. J. Fischer, "Data traffic performance of an integrated circuit- and packet-switched multiplex structure," *IEEE Transactions on Communications* **COM-28,** No. 6, 873–878 (June 1980).
28. For example, as trunk transmission costs drop it is possible in some situations for the cost of switching and control hardware in the nodes to become so significant that earlier conclusions about relative transmission costs are overshadowed by considerations of switch costs.
29. D. Boettle and M. Klein, "High speed (140 Mbit/s) switching techniques for broadband communication," *Proceedings 1986 International Zurich Seminar on Digital Communications*, March 11–13, 1986; Paper C4, IEEE Catalog No. 86CH2277-2; may be obtained from the IEEE, 345 East 47th Street, New York, NY 10017.
30. Either approach can be combined with more efficient voice compression to gain still better trunk utilization.
31. C. J. Weinstein and E. J. Hofstetter, "The tradeoff between delay and TASI advantage in a packetized speech multiplexor," *IEEE Transactions on Communications* **COM-27,** No. 11, 1716–1720 (November 1979).
32. As implemented today, burst switching has what the present study would consider a deficiency in that it supports granularity of attached line speed only down to multiples of 64 kilobits per second rather than 8 kilobits per second.
33. One mitigating factor is that message traffic necessary to continually update the topology data base is less in the centralized case, because only one such data base need be kept. For networks that are extremely volatile topologically (e.g., in case of frequent power-up and power-down of network nodes), this could become a factor to be considered.
34. R. W. Muise, T. J. Schonfeld, and G. H. Zimmermann, "Digital communications experiments in wideband packet technology," *Proceedings 1986 International Zurich Seminar on Digital Communications*, March 11–13, 1986; Paper D4, IEEE Catalog No. 86CH2277-2; may be obtained from the IEEE, 345 East 47th Street, New York, NY 10017.
35. P. O'Reilly, "Burst and fast-packet switching: performance comparisons," *Proceedings 1986 INFOCOM* (May 1986), pp. 653–666.



**Paul E. Green, Jr.** *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598.* Dr. Green is a Research Staff Member in the Computer Science Department. He received the Sc.D. degree from M.I.T. in 1953. After many years at M.I.T. Lincoln Laboratory, he joined IBM Research in 1969 as Senior Manager. His interests have centered on signal processing and computer network architecture. He has had several tours of duty on the Corporate Technical Committee and Engineering, Programming, and Technology Staff at Armonk, New York. He recently received an IBM Outstanding Innovation Award for his work in formulating and promoting Low Entry Networking. He is a Fellow of the IEEE and a member of the National Academy of Engineering.

**Dominique N. Godard** *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598.* Dr. Godard is currently a Research Staff Member in the Computer Science Department, on temporary assignment from the La Gaude Laboratory, Communications Products Division, France. He joined IBM in 1974 at La Gaude, where he was involved in signal processing and modem architecture, first as an Advanced Technology Advisory Engineer, and later as Manager of the Modem Architecture Department, Signal Converter Products. He received two IBM Outstanding Innovation Awards for his work on IBM 386X and 586X modems, and four IBM Invention Achievement Awards. He received his Ph.D. from the University of Paris, France, in 1974. He is a member of the IEEE and the European Association for Signal Processing, and is on the Editorial Board of *Signal Processing*.



Reprint Order No. G321-5286.