THE EVOLVING NETWORK

A. Jay Gross
Product Manager, Multiplexers
Paradyne Corporation
(813) 530-2785

Paper No. 0051

THE EVOLVING NETWORK

The ideal data communications network is invisible to the user and appears to be no more than a piece of wire or direct connection between the originating and destination points. And like the local electric power company, it should be so reliable and easy to use that the user takes it for granted. Today, networks have the ability to approach this ideal. The array of existing tools and emerging technology makes this possible. In many respects, the future is here today.

The beginning of modern data communications can be traced back to the 50's. It was then that American business and industry began a trend towards decentralization that created new problems and set the stage for the migration of centralized data bases out to remote locations. It became apparent that new technology would be required to transfer information from one location to another. The pressure was on for business equipment manufacturers to develop methods and systems for moving the massive volumes of information being generated. Solutions to the communications bottleneck had to be found.

The modem, which became the foundation for an entire new industry, ushered in the era of data communications in the 60's and distributed processing in the 70's. Most early concerns were easily addressed even though available options were few. Simple point-to-point and multidrop configurations predominated. Most applications required less than 2400 bps with 110 bps and 300 bps dial being used extensively into the mid and even late 70's.

The network planner's goal for providing the user a cost-effective and efficient data communications system was always met. After all, as traffic and locations grew, users could add more cheap lines and drops, along with modems that had started decreasing in price. In many cases, the best use of resources or most effective network design had very little to do with actual system implementation. Multiple lines to single locations were common.

Interestingly, at that time network control and diagnostics were considered to be of little value. By process of elimination, a problem usually could be narrowed down to one of two network components--lines or modems. Besides, no more than two points of contact for service were usually required - the telephone company or the modem vendor. Many times, one call to Telco took care of everything since both the modems and the lines were leased from them.

The good old days - less demand from users, reasonable costs, few products and services to deal with, simple problem determination and resolution. Uncomplicated network solutions that worked!

Well, the 80's changed all of that. It seems that almost overnight, users had more applications than ever, and of course, they all required better response times. Remember what happened when everyone bought a personal computer and wanted it on line.

Or, when the local telephone company increased their access charges dramatically, as well as their lead times for lines. And loss of a single point of contact for service made life unbearable as users screamed about uptime or the lack of it! Increased product offerings from multiple vendors became a double-edge sword. Often, older equipment was not compatible with newer technologies. And, of course, service offerings from the Telco's based on voice standards only added to both the confusion and frustrations. The network could grow no further using existing hardware.

The communications evolution has now reached its next logical step. The changing technology, regulations, pricing and user requirements are the causes of both today's problems and opportunities in network design. We have entered an era where the network has become as critical as the information source itself! Many businesses in our increasingly service-oriented society have discovered that the network is the tool that gives them the edge over competition. Most have discovered the time value of information.

Attention must be given to understanding the nature of the network from a business perspective - the products, markets, competition, and underlying management philosophy. Ask where did we come from, where are we today, and where are we going. Thus, as planners look toward growing the network, their goals must be to:

- decrease costs
- increase profitability
- increase productivity
- increase network availability
- seek new markets and business opportunities

Modems alone can no longer meet these goals, although they continue to play an important role in the network of the 80's. The network planner must now seek consolidation by acquiring technology that will add value in terms of both cost savings and increased function to existing systems. Thus, the trend is towards consolidation and integration (figure 1). A broad group of devices has emerged as a powerful way to meet these objectives: multiplexers, concentrators, and nodal processors.

Their primary function is to concentrate a large number of low-speed incoming lines onto one or more high-speed transmission facilities. This functionality, coupled with highlevel, built-in intelligence, not only reduces line costs but provides a uniform means of dealing with the total communications requirements of the organization. Enhanced service offerings, increased network availability, improved diagnostics, and a migration path for future applications are only a few of the benefits.

Simple, low-end concentrators include modem-sharing devices, point-to-point Statistical Time Division Multiplexers, telephone central office D4 channel banks, and point-to-point T1 multiplexers. Advanced designs range from networking statistical multiplexers and T1's to packet switches. These intelligent multiplexers are usually referred to as communications processors or nodal processors.

Unifying network architecture through the use of multiplexers takes the comprehensive approach to consolidation and integration. But, implementation of such a design must be well thought out. In addition to the business goals just mentioned, there are obviously technological issues. Everything from type and volume of traffic to applications, network control, and geography must be taken into account.

Three basic technologies have emerged with capabilities that can meet the needs of different situations. Statistical multiplexers, packet switching, and T1 are the choices to be considered.

Statistical multiplexing is used primarily in situations where asynchronous communications predominate. Because this method dynamically allocates bandwidth as needed, the short character-at-a-time transmission used in asynchronous communications can produce tremendous efficiencies. The ratio of aggregate channel speeds to the link speed can easily be 10-to-1 or greater. In other words, ten devices of 4800 bps each, for a total aggregate 48,000 bps, could be statistically multiplexed over a link running at 4800 bps (figure 2). Not only is the cost of nine lines saved, but the speed of the modems used between the two multiplexers is much less than one might expect!

This is not to say that the laws of physics are altered. Instead, a statistical multiplexer takes advantage of the fact that all of these devices probably won't be transmitting at the exact same moment in time. Even in a "heads down" order entry application, an asynchronous terminal is sending data only about 10 or 15 percent of the time during an eight-hour day. Besides, how many operators do you know that can type at 4800 bps? That would be the equivalent of about 36,000 words per minute!

To look at it another way, if each operator were typing at 40 words per minute, which is about 5 bps, 900 devices transmitting simultaneously would be needed to fill up 4800 bps worth of bandwidth! Of course, there are other considerations which, from a practical point of view, make this compaction ratio much higher than we could really obtain. The point is, all users are virtually guaranteed a time slot on demand. That's why the apparent throughput in this example is 48,000 bps even though the real rate of the link is only 4800 bps.

Here's how its done. The data from each device forms a frame that is sent to the remote statistical multiplexer. Included with this frame is additional information for error control, addressing, flow control and signaling. Most multiplexers use an international standard protocol for framing called HDLC (high-level data link control). This is the same transport protocol used in X.25 packet switching networks (figure 3).

Each frame can contain up to 256 bytes of information, although most statistical multiplexer vendors use 128 bytes. Using this technique, all users will get at least several bytes of data into a given frame if requested. If there is not enough data to fill up the 128 bytes, the frame will automatically adjust down. In our example, the ten terminals will have at most 40 bytes per minute times 10 devices, divided by 60, or a total of about seven bytes of data to send per frame.

As efficient as this is, another 20% can be gained by stripping the start, stop, and parity bits of the asynchronous data at the sending statistical multiplexer and then adding them back at the receiving end.

Another important feature of a statistical multiplexer is its ability to buffer data. Should all operators manage to hit a key at the exact same moment in time, data will not be lost. A buffer size of 4 to 16K will usually suffice, depending on the number of terminals and their respective speeds. This situation will change as printers and batch devices are added. Compaction ratios of 10-to-1 are no longer valid. Instead, a more adequate rule of thumb would be ratios of 4-to-1 for printers (up to a few pages at a time) and 2-to-1 for batch devices.

Instead of increasing the buffer size in order to handle the demand created by these bandwidth hogs, flow control is used. When about 80% of the buffer is full, a command is issued by the statistical multiplexer to both the local DTE's and the remote statistical multiplexer. The remote multiplexer will, in turn, signal the sending device to stop sending data. The most common method of doing this is an in-band scheme known as Xon/Xoff. As the name implies, the proper software flow control signal is sent to start and stop normal data flow from the originating device to its destination. Xon will usually be given when the buffer empties to about 40% of its capacity. Because this activity is in-band and automatic, the operator will have no knowledge of the occurrence.

Hardware flow control can also be used. EIA control signals such as Clear to Send (CTS) and Data Set Ready (DSR) are used.

As most users already know, there is not much in the way of error control in the asynchronous world. At least not until statistical multiplexers arrived on the scene. Like most synchronous protocols, HDLC uses a redundancy check algorithm in conjunction with an Automatic Request for Repeat (ARQ) for end-to-end error control. So, not only is every device virtually guaranteed a time slot, users can count on the information arriving at its destination error free!

Today's network planner can utilize the intelligence inherent in a statistical multiplexer to take this one step further. A comprehensive network solution can be implemented by using building blocks available from one of a handful of networking statistical multiplexer vendors. From four channels to 240 channels, an integrated family of products will allow an asynchronous network of any size to be designed. Some vendors have even integrated multiplexers into their network management and control centers. Now a single point of control exists for multiplexers, modems, and DDS equipment such as DSU/CSU's (figure 4). The ultimate approach for consolidation and control of the asynchronous network would certainly seem to be statistical multiplexing.

What happens when we consider synchronous data? In a mostly asynchronous environment with some synchronous traffic, today's statistical multiplexers are the way to go. But, as the percentage of synchronous traffic approaches 20%, alternatives should be considered. And, when this number exceeds 40%, the statistical multiplexer becomes inadequate very quickly. Packet switching may be a better answer.

A packet switch is no more than a very sophisticated statistical multiplexer. As stated earlier, both use HDLC as their protocol which is considered to be the transport layer of X.25, known as X.25 level II. Packet switching utilizes the next layer which is X.25 level III. In addition to the HDLC protocol, this level contains an advanced addressing structure resulting in only one channel being assigned to one frame. The result is block oriented (figure 5).

This multilayered addressing structure, combined with the enormous power built into packet switches, gives it the ability to support diverse vendor-specific protocols. Thus, packet switching seems to better satisfy synchronous network design applications, particularly large networks with multiple synchronous protocols.

Packet switching also differs from statistical multiplexing in its hardware makeup. There are packet nodes and packet assemblers/disassemblers (PADs).

The node forms the network backbone. Its software enables the node to communicate with adjoining nodes so that traffic information is constantly exchanged and updated. In other words, it acts as a trunk-interfacing tandem switch. Like a telephone company central office, call requests are properly directed to the destination node over the path with the shortest delay.

The PAD connects remote terminals directly into the network nodes. Typically, they will packetize several devices of a specific protocol into a single X.25 network interface. The PAD can be located locally to the node or remotely via a pair of modems and a dedicated line. At the host site, a PAD can be installed so that, in effect, the packet network becomes transparent by interfacing at the port level, or instead, one of the node trunks could support a connection to a single host port configured to de-mux multiple virtual circuits at the X.25 level (figure 6).

Another important feature of packet switching is the sophisticated network management tools available. In addition to network control and diagnostics, a packet network control center can provide capacity management as well as usage accounting and billing.

Capacity management ensures that packets are sent to the receiving party by the most efficient available route, based on network traffic at the time of call setup. This operation automatically smooths peak load that otherwise would require additional capacity.

Usage information can be important to organizations that charge network time to departmental cost centers or customers. The advantages of packet switching are becoming clear. Its ability to provide a vendor independent solution for large diverse networks is the key. Packet thrives on multiple hosts, protocols, applications, and locations.

The third area of multiplexer or concentrator technology is T1.

In recent years, the term T1 has been used to refer to any digital arrangement operating at 1.544 Mbps. T1 has been used in this country for almost 25 years, yet for many, it continues to be a vague and mysterious technology. And, indications are that T1's emergence as the transport technology of the future is only now beginning.

The first T-carrier facility was introduced into the otherwise analog telephone network in the early 60's to interconnect central offices. By digitizing voice signals and multiplexing them using time-division techniques, the T1 link permitted 24 voice-frequency channels to be carried over just two pairs of wires. The first generation of T1 central office termination equipment was the D1 channel bank.

By the late 70's, this evolved into the D4 channel bank. Although voice digitizing and channelling techniques had changed, basic technology remained the same; a very large time division multiplexer (TDM) with 24 channels and a link speed of 1.544 Mbps (figure 7a).

In 1984, several things happened to change the way T1 was perceived and utilized. First, AT&T made T1 more accessible by increasing the number of facilities installed and lowering the cost. This became known as Accunet 1.5. With the advent of divestiture, other carriers, including the Regional Bell Operating Companies (RBOCs) increased competition and availability by installing even more high-speed circuits, much of it being fiber.

It was at this point that vendors began to see the real potential T1 offered for voice and data integration. An updated version of the D4 channel bank was introduced--second generation equipment simply known as T1 multiplexers. By adding microprocessors for intelligence, the bandwidth could be subdivided into more than 24 channels. Supervisory capabilities were added and limited networking was achieved (figure 7b). At about this same time, the idea of a voice and data integrated services digital network (ISDN) started to emerge with T1 playing a key role.

What seemed like a stable and known technology with interface and framing specifications cast in concrete suddenly became complex and overwhelming primarily due to a lack of understanding in the data communication marketplace. Even a fundamental grasp of these concepts will lead to advanced networks able to take advantage of T1 economies of scale and newer equipment designs.

What does it mean to be T1 compatible? Today, the answer is not a simple one, as there are different levels and framing formats. Obviously, to start with, the multiplexer must at least transmit and receive a signal at 1.544 Mbps. This is DS-1 compatibility and includes electrical characteristics.

When a DS-1 signal is used for D4 service, it consists of frames of 193 bits each. This represents an 8-bit byte for each of the 24 sub-channels, plus an extra bit for framing. Sampled at 8,000 times/second, each of these 24 slots represents 64 Kbps of bandwidth or a total of 1.536 Mbps of usable bandwidth. The 193rd bit, also sampled at 8,000 times a second, accounts for the remaining 8 Kbps of the total 1.544 Mbps bandwidth. Twelve of these 193 bit frames are known as D4 superframes and represent the framing level of compatibility. (See Table 1)

The next level is D4 Channelization, required of T1 circuits that terminate at the central office. A DS-1 signal made up of 24 subchannels, each taking 8 bits at a time, is now needed. These are numbered DS-0 1 through DS-0 24. For premises-to-premises transmission, the user must only maintain the 193rd bit pattern--everything else is transparent.

The third and final level is signaling compatibility. Here, the 6th and 12th frame of every superframe have the 8th bit of the user's information robbed. These signaling bits are used for basic telephone control, such as on-hook/off-hook indications.

It is only when each of these three levels of compatibility are met that a T1 multiplexer is truly D4 compatible; framing, channelization, and signaling. It is at this point that the user can take full advantage of new and existing tariff offerings.

Since they are TDM's, T1 multiplexers can be bit or byte-interleaved. Bit-interleaving, being the most efficient of the two, involves transmitting each bit as it is received from the incoming channels. The advantage is that, unlike statistical multiplexers and packet switching, only a small amount of buffer is needed--on the order of several hundred bytes per port or less. Since the multiplexer does not need to wait for an entire character to arrive before putting information on the T1 pipe, bandwidth is maximized and delay is minimized. However, bit-interleaved multiplexers cannot be used to interface with many of AT&T's

service offerings such as Customer Controllable Reconfiguration (CCR), also known as DACS (digital access and cross-connect system). The ideal T1 multiplexer is programable as either bit or byte-interleaved. As a practical matter, bit-interleaving can and should be used in most situations.

Today's T1 multiplexer has become so powerful that it is correctly referred to as a nodal processor, supporting multiple links, large channel capacities, and features such as automatic alternate routing and bandwidth contention. Virtually any application can be supported, including synchronous, asynchronous, voice, video, and high-speed data needs right up to 1.536 Mbps. Integrating these applications appears to be the key to using T1 effectively (figure 8).

Table 2 compares the relative merits of each of the three technologies discussed.

CONCLUSION

Because there are choices, it is tempting to pick the one direction that seems to offer the best overall fit. But, do network planners have to lock themselves into one way of thinking or one type of technology? Most corporate networks are really a collection of different applications and even separate networks based on common business interests. Thus, it may make sense to examine the integration of these technologies.

Integration can take place by simply sharing high-speed trunks to common locations. This allows smaller, independent networks to remain autonomous while reducing the overall line costs. But, as each individual network grows, reallocation of the backbone trunk bandwidth will be necessary, resulting in performance constraints. Also, sharing is fixed to the extent that traffic on one network cannot borrow bandwidth from the others during peak traffic periods or in times of link failures.

A better and more fully integrated approach is the hybrid design. This network also uses high-speed backbones, but combines private and public facilities as needed. This mix and match approach tailors the architecture to traffic characteristics.

The mechanism for interconnections are multifunction gateways. More than protocol converters, gateways are intelligent access points that integrate one network into another. As applications and requirements change, so does the personality of the network. Gateways enable the network planner to take advantage of existing service offerings from AT&T. Currently, these include CCR, M24, and Software Defined Network (SDN). Public and private packet networks can also be accessed.

Then there's ISDN. Touted as the ultimate solution for integrating all forms of information, ISDN is being carefully designed to retain compatibility with existing switching and transmission equipment, most notably T1. In fact, the number of rapidly growing T1 based private backbone networks are in effect private ISDN facilities. Thus, a network solution incorporating T1 will allow migration towards ISDN as it becomes available.

The best solution for network needs is based on hard questions and even harder decisions. Those willing to invest the time and energy in this process will develop a comprehensive communications strategy for today <u>and</u> the future.

POINT-TO-POINT LINES

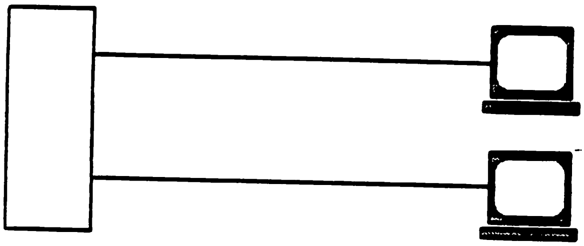

paradyne

0051-10

CONSOLIDATED NETWORK ?

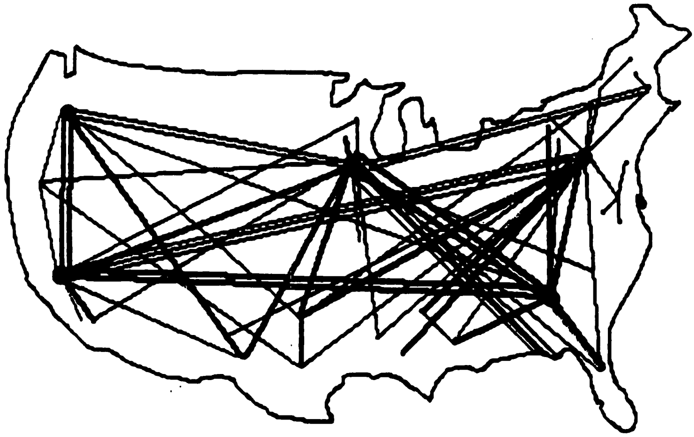

paradyne

0051-11

STATISTICAL MULTIPLEXING BACKBONE

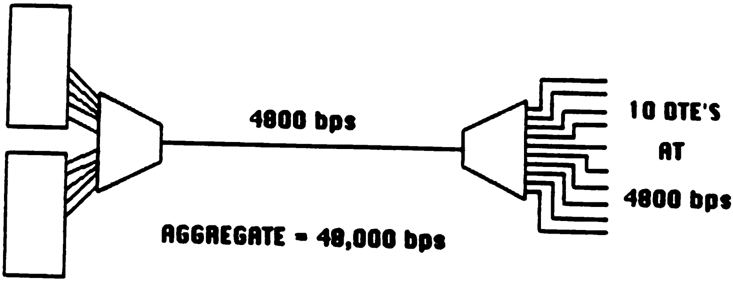

paradyne

0051-12

PACKET SWITCHING NETWORK

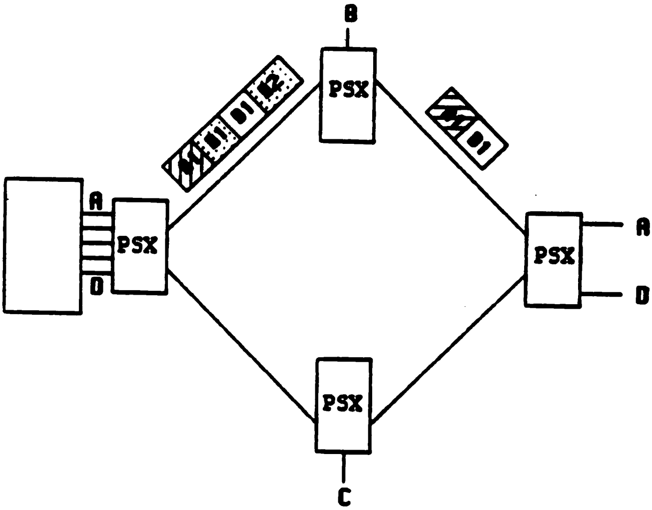

paradyne

0051-13

T1 TECHNOLOGY

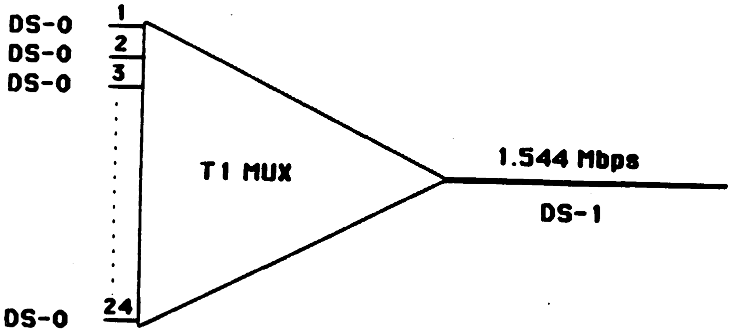

8 BITS X 24 CHANNELS = 192 BITS

192 BITS X 8,000 SAMPLES/SEC = 1.536 Mbps

8 BITS X 8,000 SAMPLES/SEC = 64 Kbps

paradyne

TABLE 1

STATISTICAL MULTIPLEXING
Vs.
PACKET SWITCHING

	STAT MUX	PACKET
ALTERNATE ROUTING	GOOD	EXCELLENT
ASYNCHRONOUS	EXCELLENT	GOOD
SYNCHRONOUS	FAIR	EXCELLENT
INTERACTIVE	EXCELLENT	EXCELLENT
BATCH	FAIR	FAIR
ERROR CONTROL	EXCELLENT	EXCELLENT
MANAGEMENT	GOOD	EXCELLENT
HIGH SPEED/VOLUME	FAIR	FAIR
COST	LOW	HIGH

paradyne

THE HYBRID NETWORK

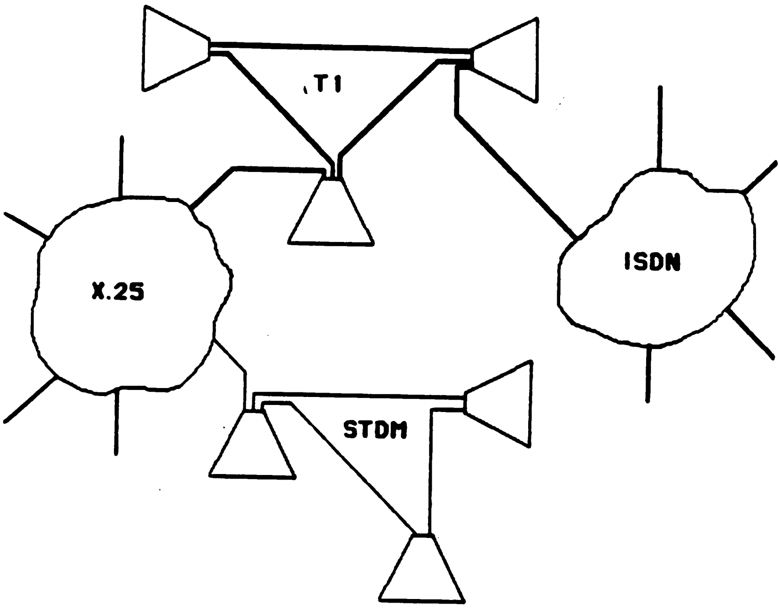

0051-16