Metropolitan Area Networks: Where Many Standards Meet

James F. Mollenauer
Chair, IEEE Project 802.6, Metropolitan Area Networks

Computervision Corporation
Bedford, MA 01730

Abstract

Metropolitan area network (MAN) standards represent an extension of local area networks to citywide dimensions. Because of the economies of scale in integrated circuits, standards will be necessary before the field develops fully. MAN's must be capable of interfacing to local area networks on one side and wide area networks on the other, while being constructed out of telephone company fiber optic links internally. Development of MAN standards requires consideration of standards in all of these areas. Two proposals currently active in the IEEE 802.6 committee will be described.

Introduction

The area of standards for Metropolitan Area Networks (MAN's) has recently become a very active one. Many currents in standards and their underlying technologies are coming together in a way which should offer considerable benefits in the future.

Metropolitan area networks are an outgrowth of local area networks, an area that has received a great deal of attention for the last five years. While the IEEE Project 802 local area networks are limited in their range to a set of nearby buildings, the technologies being studied in 802.6, the metropolitan area networks group, are capable of spanning a whole metropolitan area.

The MAN effort was started in 1982, about 2 years after Project 802 was initiated. The group's charter was to provide standards for networks optimized for areas of 50 kilometers in diameter. This did not exclude technologies good for 200 kilometer runs, or prohibit installations smaller than 50 kilometers in size. It did distinguish MAN issues from those of wide area networks, where the protocols are expected to be distance-insensitive.

From the beginning, it was assumed that broadband (CATV), fiber optics, and radio were valid choices of media. The functional goals were to provide for interconnection of LAN's, bulk data transfer, digitized voice and video, as well as conventional terminal traffic. In meeting these goals, the 802.6 committee has been free to make its own choice of media-access protocols and physical parameters.

Initially 802.6 concentrated on standards applicable to CATV-based systems, due to the availability of bandwidth and spare cables in the CATV franchises. However, in the last three years attention has shifted to fiber optic facilities. In deference to the speed characteristics of fiber optics and the capacity needs of backbone systems interconnecting LAN's and other large data sources, the IEEE Standards Board in 1985 approved lifting the 20 megabit/second maximum speed for Project 802 LAN's. The MAN can go as fast as the fiber (or other media) permits.

In the case of fiber-based networks, the obvious choice of operator is the telephone company. Large quantities of fiber have been installed by the telephone industry in the last few years, and as with the CATV systems, it is desirable to aim standards toward compatibility with the installed cabling. There is no need to build MAN's from scratch if suitable media are already in place.

We are attempting to draft standards for MAN's that operate in conformance with the principles of IEEE 802 local area networks, but are constructed of links that must also retain compatibility with other standards such as ISDN. This brings us into contact with a very broad cross section of standards activity.

Pro-Active Standards

A role for official standards in establishing the cutting edge of technology is relatively new. Usually standards have been used to codify design practice and to impose some uniformity on a technology as it matures. This has changed at least in the data communications area. Here the effect of some standards has been to define the state of the art and to provide a new point of stability where users can build networks that will not become obsolete in a short time.

The best example of this is X.25 for packet networks. X.25 was standardized in CCITT before anyone had implemented it, in 1976. A number of network operators (mostly in Europe) felt a need to define a standard quickly, and

CH2549-4/88/0000/0002$01.00 © 1988 IEEE

the operating experience with the Arpanet system in the U.S. was distilled into an official standard. Of course there were deficiencies, and subsequent revision in 1980 was required to correct them and provide a basis for most of the commercial packet networks.

Despite the need for revisions, the example of X.25 has proven infectious. The economies of scale that are possible in silicon make standards extremely attractive. The cost of designing and producing integrated circuits is almost independent of the number produced. Very complex and highly general data communications facilities can be produced at low cost if the volumes are high. The existence of standards provides assurance of high volumes and attracts IC manufacturers. Otherwise no vendor is willing to go beyond the lower development investment of low-scale integration or gate-array implementations. The VLSI solution, when volume production is possible, is very much cheaper.

The IEEE 802 family of standards likewise illustrates the pro-active role of standards. While Ethernet had existed for several years as a defacto standard published by Xerox, Intel and DEC, it was not until the 802.3 standard was firm that VLSI chips supporting Ethernet appeared on the market. In the case of the 802.4 token-passing bus, no equivalent system had been marketed before the standard was approved; significant numbers of installations are only now appearing, using the VLSI chips that have recently been manufactured. The perils of working close to the edge of technology are also illustrated. Some parts of the 802.4 standard had to be reworked because of difficulties in building production networks according to the standards.

In the 802.6 arena, the pro-active role of standards is prominent. The proposal the committee is working on, known as QPSX or Dual Bus, is currently being tested in field trials in Australia. Widespread deployment is awaiting VLSI silicon, which in turn depends on the existence of standards for assurance of sufficient market beyond Australia. Similar considerations have also applied in the case of the MST proposal, which depended on a packet media-access layer derived from work in other standards committees.

MAN's versus LAN's

A metropolitan area network is an outgrowth of the idea of a local area network, which in turn is a descendent of a computer's internal bus. Almost all computers are based on an internal bus; if that bus is serialized for greater distance range and the access allocation is distributed to each unit, then we have a local area network. Its great utility derives from the fact that it still runs fast (delivering a bit rather than a whole word per clock cycle) and is well adapted to computer traffic, which is very bursty. This contrasts with data communication technology based on the telephone system, which is well adapted to a continuous stream of data at a constant rate.

In the case of the metropolitan area network, the optimization distance is raised from several kilometers (the limit of Ethernet) to 50 kilometers in diameter, the size of a whole city and its suburbs. This scale originally derived from the needs of the satellite communications industry, which needed low-cost but high-speed links from ground stations to the customers' buildings. In addition to data, this traffic had a high proportion of digital voice and occasionally video as well.

This was the genesis of 802.6. As it has turned out, the importance of satellite communications has waned under the impact of fiber optics. However, the need for high-speed interconnection of the LAN's that corporations have installed in their various sites has increased with time. The view of the MAN as a "last mile" in a satellite system has been replaced with a more earthbound one: the MAN can interconnect corporate premises in the same city, and perhaps with different optimization can also serve small businesses and home computers. Initial work was based on CATV technology, but in the last three years the work has shifted to fiber optics. The great interest in fiber for other communications applications has been matched in the proposals submitted to the MAN standards committee.

The utility of standards for distances intermediate between LAN's and WAN's (wide-area networks) results from two factors, one in the market area and one technical. From a marketing standpoint, there are many organizations that tend to occupy buildings throughout a city and its suburbs as they grow. Whether the buildings are adjacent on a campus or scattered over distances of miles, they have a need for volumes of communication between them. Because of the distances involved, and soon because of speed requirements as well, the present IEEE 802 networks are inadequate; different technology must be employed.

On the technical side, the emergence of fiber optics has provided both higher speed by a factor of 10 to 50 and relief from the distance limitations of LAN's. The distances that can be served with shared-media access methods are still not unlimited, but they are definitely appropriate for metropolitan dimensions. Hence the MAN is a viable technology different from both local and wide area networks.

Current Proposals

After a year of study, the 802.6 committee has voted to move ahead with a standard based on the Dual Bus, or QPSX, proposal. This proposal is very distinct from any of the other protocols used in standard networks. It was introduced by Telecom Australia as an evolution of a coax-based system that originated at Bell Laboratories. QPSX (for Queued Packet and Synchronous Exchange) uses a dual fiber (or other medium) formed into a physical ring, with each fiber running at 150 megabits per second. It is not a true ring, however, because one unit does not repeat the data. This is shown schematically in Figure 1. Thus it is properly a bi-directional bus, and packets just "fall off the end" rather than being explicitly deleted from the ring as in the token ring protocol. This feature confers a great deal of robustness: should a cable segment or node fail, the logical break in the ring is moved to the location of the failure, and the system continues to operate with no degradation in performance or propagation time. This is shown in Figure 2.

Also, the bidirectional nature of the bus makes it possible to use a very efficient scheduling method. Units waiting to transmit to the right, for example, count down from the number of unsatisfied slot requests originating to their right, arriving on the other bus. They then decrement the count by one for each empty slot observed going toward the right. When the count reaches zero, transmission is enabled for the next empty slot on the right-bound bus. This is shown in Figure 3. This distributed scheduling algorithm guarantees first-come-first-served access to the network, with no loss of efficiency under high load.

A further benefit of the dual bus technology is that all data can be ORed to the bus. The number of electronic components in series with the data stream is very small, resulting in higher reliability and lower cost.

The alternate proposal that was studied in the committee is MST, or multiplexed slot and token. The packet handling capability of MST is based on an underlying token protocol for media access. It overlays this with the capability to provide isochronous, or fixed bandwidth, channels for applications that require it. MST has several speed options to match the speed characteristics of the telephone links that will comprise the MAN's. MST normally runs in packet mode but on command it shifts into the isochronous mode, suspending packet transmission. When the fixed-bandwidth transmission is done, it shifts back to packet mode.

<u>AT&T</u>: <u>Interconnecting</u> <u>the</u> <u>MAN</u> <u>Loops</u>

A second item of work accepted by 802.6 is a proposal by AT&T Bell Laboratories for interconnecting multiple Dual Bus loops with a multi-port bridge operating at the MAC level. Many of the Project 802 working groups are now grappling with the problems of retrofitting MAC-level bridging into their architectures. In 802.6 we intend to make provision for such bridging from the beginning, and to take advantage of the latest technology to construct these bridges.

The multi-port bridge is in fact a small but very fast switch. The telephone industry is now developing packet switches that are capable of operating at fiber-optic speeds. By using intelligent crosspoints, these switches are able to interpret the headers of small packets on the fly and route them through a "banyan network" named for a tree with a complex branching structure.

Such switches can serve to interconnect multiple MAN's into a coordinated system. While the 802.6 MAN is capable of transmitting 300 megabits per second, this does not exhaust the future metropolitan network needs. Partition-

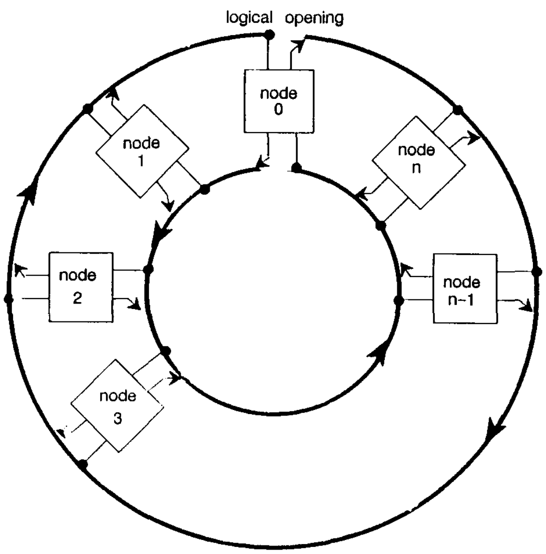

1. QPSX in normal mode. The controller acts as the logical end of both busses, preventing the flow of data past it. Writing to each bus is unidirectional as indicated by the arrows.

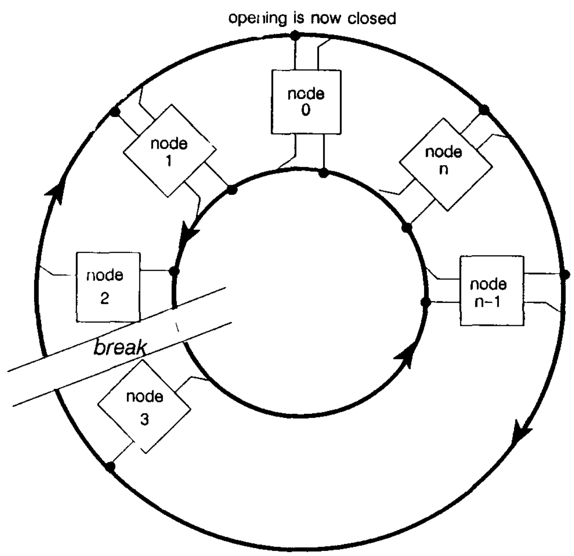

2. QPSX recovery. When a failure occurs, the logical opening is shifted to the location of the failure, and the controller now passes data through in both directions. Network performance is not degraded.

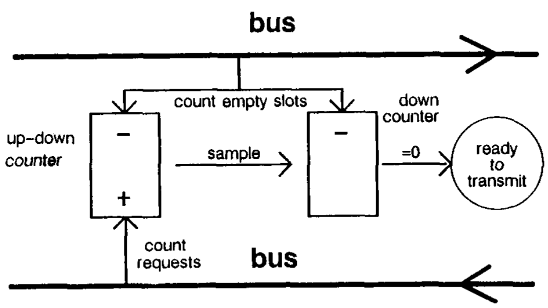

3. QPSX scheduling. See text for explanation.

ing a metropolis into several networks even when they are well short of capacity is likely to provide better reliability and maintainability. The multiport bridge can then interconnect the MAN loops with hardly any limit on the total capacity of the system.

Interaction with Other Standards

Because they sit in an intermediate position, MAN's need to interact with both local area networks and wide area networks. Compatibility with each side has advantages but confers some restraints.

If a MAN carries voice, calls originating on the MAN will often terminate outside it. Continuity signaling for call setup and control purposes must be established, just as it is today for routing calls through a sequence of switching centers and operating organizations. In 802.6 we do not plan to reinvent this process; we will use the CCITT Q.931 protocols for voice call handling. Provision will be made for call-control processors that will implement these protocols in a way that interfaces cleanly to the MAN.

When installed as a private network analogous to a leased telephone line, a MAN does not have severe security problems. Only one company's data is on the fiber, and there is no problem inherent in bringing that fiber in and out of the company's buildings to connect to LAN's, computers, and PBX's. However, the bandwidth is well beyond the requirements of many organizations, and there will clearly be economies in sharing the fiber. In this case, the MAN becomes a public facility to be used by many people.

This poses problems of billing, which are not hard to solve, and of security, which are much more difficult. There is a strong feeling that it is very undesirable to bring a public fiber into user premises. Wires or fibers entering a building should contain only that user's information. This means that a link must be made between the user's facilities in the building with the network outside. A separate access network is required, with a gateway or bridge to segregate the public and private portions of the network. This is shown in Figure 4.

The access network may be either a private installation of an 802.6 network, an 802 LAN, or a point-to-point link of the type being standardized by the 802.9 Integrated Voice/Data working group. We will need to provide for gateways to handle all these possibilities, and in addition many vendors will probably support access networks with proprietary protocols like Decnet.

The flexibility of the Dual Bus system to use fiber or coax, or indeed even radio, stems from the use of the CCITT G.703 interface. All stations must use this interface; behind it the installation can include transducers to convert the station output to optical signals, or it can be run on coaxial cable.

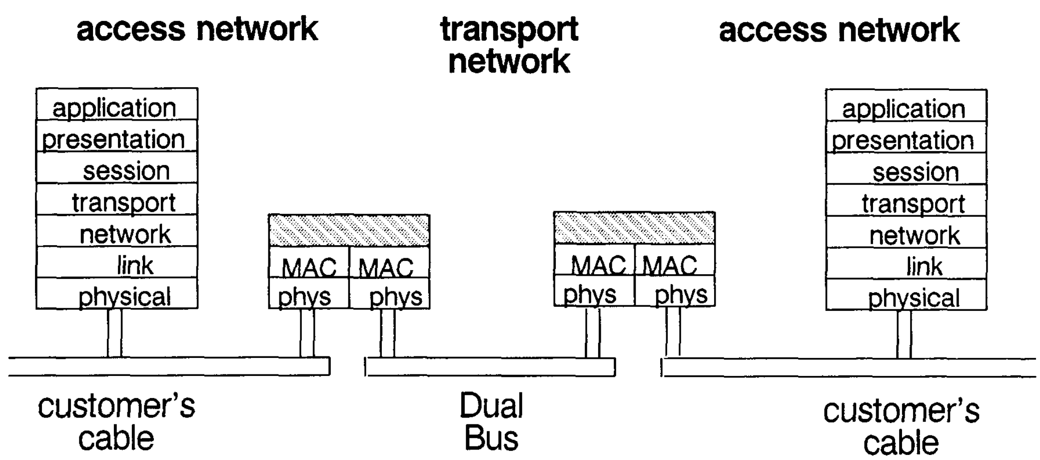

4. MAN architecture. When a cable is shared by multiple users, it does not pass through each user's premises; data for that site is taken from the transport network by a MAC-level bridge and passed to the access network.

Another standard in the fiber-optic area is FDDI, or Fiber Distributed Data Interface, from American Standards Committee X3T9.5. FDDI was proposed as a computer-room interconnect, but it is capable of serving as a high-speed LAN with a total ring length running to 200 kilometers or more. Originally FDDI was considered for MST's packet media access layer. However, FDDI is fixed at 100 megabits per second, which is not optimal for networks constructed out of telephone links. As a result, work shifted to a consideration of 802.5's newly-proposed early token release option, which can greatly improve the efficiency of token rings at the speeds and dimensions of MAN's.

Meanwhile, X3T9.5 has also been looking at metropolitan networks. Given that the distance limits of FDDI are appropriate, fixed-bandwidth capability must be added for voice and video. This is done in the FDDI-II proposal by dividing the 100-megabit bandwidth of FDDI into 16 channels of 6.144 megabits each. These channels must be further subdivided into voice channels. The fixed 100-magabit speed precludes the use of standard telephone circuits; as a result, FDDI-II may find its best application in campus networks using private cables.

The links that compose the MAN will probably be standard circuits used in the telephone industry. While it is possible to provide users with "dark fibers" and leave all operation in the user's hands, there are many benefits in using the standard links. Installation, maintenance, and automatic backup can all be done more easily with the normal fiber circuits. In the case of public fibers this is particularly important, since the utility must be able to measure the traffic in order to bill for it. The present DS3 at 45 megabits and the future broadband ISDN at about 150 megabits are examples of standard links that would serve well as a MAN basis.

In order to provide consistency with broadband ISDN, 802.6 has recently established a close liaison with T1D1.1, the main American committee for broadband ISDN issues. The 150-megabit speed of broadband ISDN is ideal for digital video, but few other applications are able to supply or absorb data at these rates on a full-time basis. Metropolitan area networks are able to provide an effective multiplexing arrangement for BISDN, while BISDN provides the underlying transport for the MAN.

We expect to be able to exploit this synergy in drafting the MAN standard and in the installations that will bring metropolitan area networks into operation.