Introduction to 100VG-AnyLAN and the IEEE 802.12 Local Area Network Standard

100VG-AnyLAN provides a 100-Mbit/s data rate with guaranteed bandwidth and maximum access delay for time-critical applications such as multimedia, using existing building wiring. It uses demand priority protocol. Developed by Hewlett-Packard and now supported by over 30 companies ranging from integrated circuit vendors to systems suppliers, demand priority is well on its way to becoming the IEEE 802.12 standard.

by Alan R. Albrecht and Patricia A. Thaler

100VG-AnyLAN is a new, high-speed addition to HP's AdvanceStack local area network (LAN) product group. It is an economically effective upgrade path for congested 10-Mbit/s 10Base-T Ethernet and 4/16-Mbit/s token ring networks. It provides a 100-Mbit/s data rate with guaranteed bandwidth and bounded access delay for time-critical applications, using existing building wiring. This provides high performance for traditional data transfer applications. It also provides emerging multimedia applications, such as interactive video, with the low delays they require. It delivers this performance over the most common networking medium, 4-pair unshielded twisted-pair (UTP) telephone wire.

100VG-AnyLAN uses the demand priority protocol, which was developed as a joint effort by Hewlett-Packard Laboratories in Bristol, England and the HP Roseville Networks Division in California. Now supported by over 30 companies ranging from integrated circuit vendors to systems suppliers, demand priority is well on its way to becoming the IEEE 802.12 standard.

The 100VG-AnyLAN articles in this issue look at the development of the demand priority protocol and the 100VG-AnyLAN product set.

Local Area Network Technology

Before the initial development of local area networking in the late 1970s, the telephone system was the only generally available data communications option. Bandwidth (3 kHz on a voice-grade line) was clearly a problem and a number of different types of new data networks were proposed. Two, Ethernet[1] and token ring,[2] have emerged to dominate the local area networking market.

Ethernet (IEEE 802.3). Ethernet was developed in the late 1970s as a 10-Mbit/s answer to the limitations of the telephone system. It became an IEEE standard in 1985. All nodes were connected to a single central coax bus, which proved to be somewhat inflexible as users change locations or are added to the network.

The Ethernet access policy is CSMA/CD, which stands for carrier sense, multiple access with collision detection. It allows any node to transmit a packet (with up to 1500 bytes of data) anytime it detects silence (no signal) on the network. This can lead to packet collisions if two or more nodes need to transmit and detect silence at the same time. Each involved node is required to back off (cease transmitting) immediately after a collision is detected, but time is consumed and the available bandwidth is effectively reduced during high-traffic periods.

The protocol also requires each node to monitor the network traffic and to decode (filter) the destination address of each packet to determine whether it should be received by the node. Packets with the node's individual or group address are copied into memory and packets with nonmatching addresses are ignored.

The 10Base-T star topology was proposed by Hewlett-Packard in 1987 and became part of the standard in 1990. The center of the star is a network concentrator (hub) which is typically located in a wiring closet. Each node is connected to the hub by voice-grade twisted-pair cable. 10Base-T retains the basic features and access policy of the bus network and also adds a level of fault tolerance. Link faults at individual nodes are isolated by the hub and do not take down the entire network. 10Base-T has become the most common IEEE 802.3 network configuration.

Token Ring (IEEE 802.5). Token ring was proposed as a 4/16-Mbit/s solution to the Ethernet collision problem and became an IEEE standard in late 1984. The original network structure is a ring around which both tokens and information packets (up to 4500 data bytes) are passed. The network medium is IBM type 1 shielded twisted-pair (STP) cable. Token ring networks are also now commonly installed in star configurations.

The token ring access policy is designed to be both collision-free and priority-based. It prevents any node that does not currently "own" the token from transmitting a data packet, and it provides eight priority levels to allow some classes of data to take precedence over other classes.

All data packets and tokens contain an access control field that allows the successive nodes on the network both to reserve the token and to indicate their reservation priority level. The node that currently owns the token transmits its data packet with the reservation bits in the access control field set to minimum priority. Each successive node forwards the packet as it is being received. It also interrogates the destination address field to determine whether it should copy the data frame and the access control field to determine the current reservation level. If the node needs to send a data packet and the reserved priority level is less than the node's level, the node indicates its need by changing the value of the reservation bits in the forwarded packet.

The sending node removes the packet from the network and transmits a new token with the priority bits of the access control field set to the priority level indicated in the returned packet. The token then circulates to the node that first reserved that priority. That node removes the token and transmits a data packet. The token circulates continuously at minimum priority in an idle network.

The Local Area Networking Market

International Data Corporation (IDC), a market research firm, reports that the worldwide installed Ethernet base at the end of 1993 was 26,376,000 nodes, up 102% from 1992.[3] They predict an installed base of over 75,000,000 by the end of 1995, predominantly 10Base-T.

The token ring worldwide installed base was 6,744,000 at the end of 1993, up 115% from 1992.[3] IDC predicts that the token ring installed base will approach 14,900,000 by the end of 1995.

Current Network Pressures

The last ten years has seen a hundredfold increase in the speed of computers and the size of files created by sophisticated applications. Meanwhile, the data transfer rate of most networks has remained constant at 10 Mbits/s to 16 Mbits/s.

The first signs of network strain are users complaining that performance is falling off and response times are rising. The cause is almost always congestion—too many users, too much information. Network-connected high-performance desktop systems are intended to give users instant access to any appropriate information, anywhere in the organization, at any time, and this generates high levels of traffic.

Sometimes, network bottlenecks arise from individual applications exceeding the bandwidth of the network. Data-intensive applications, such as database access, image analysis, desktop publishing, network printing, and CAD, require that very large amounts of information be transferred in a single burst. For example, a desktop publishing application might require 10 megabytes for a single page incorporating several typefaces, a bitmapped logo, and four-color graphics. On a typical Ethernet or token ring network, it could take as long as 20 seconds to retrieve that one page. A multipage document could take a minute or two to retrieve.

Cable Types

Cables can be categorized in various ways: according to their physical structure, the material used for transmitting signals, and the uses for which they are suitable. Common types are listed below.

UTP: Unshielded twisted-pair, 100-ohm balanced cable. The lack of shielding makes UTP cable very low-cost, but introduces problems of cross talk when the pairs are in close proximity.

Category 3: Voice-grade cable, such as telephone wire, with 16-MHz bandwidth, used in 4-pair groups for each link. 25-pair bundles of Category 3 UTP are common in existing LANs. This is an important consideration when designing a network protocol.

Category 4: 20-MHz bandwidth, used in 4-pair groups for each link.

Category 5: Data-grade cable with 100-MHz bandwidth, used in 2-pair or 4-pair groups for each link.

Optical-Fiber: Cable consisting of a minimum of two strands of optical fiber running parallel within a protective jacket. Each fiber is usually composed of glass 125 μm in diameter, and has a 62.5-μm core. Transmission is by light beam at 850-nm or 1330-nm wavelength.

STP: 150-ohm balanced shielded twisted-pair cable. Usually used in 2-pair groups for each link.

The situation will rapidly worsen with the accelerated development of time-sensitive multimedia applications. Real-time audio and video for video conferencing and interactive video require that packets of data be transferred continuously with minimal delay. They cannot afford to have any packet delayed or dropped because of a collision or congestion on the network.

Design Goals

The problem presented to HP's network design team contained several major goals and considerations:

- Speed. The current networks are clearly too slow. A major improvement would be to increase the speed of the network. 100 Mbits/s would allow ten times the amount of traffic.
- Guaranteed Access. While multimedia and other time-sensitive applications require the increased total bandwidth that 100 Mbits/s would provide, they also need guarantees that information will get through within a stated delay window, whatever the traffic on the network.
- Cost. Existing networks have already required major investments in the wiring structure. The new network should be able to operate over generic twisted-pair building wiring. Fiber-optic cabling should also be allowed.
- Topology. To be compatible with existing wiring, the new network must use a star topology. The allowed network diameter should be at least 2.5 km with three or more levels of hub cascading.
- Software Compatibility. The network should be compatible with both Ethernet and token ring frame formats and should preserve existing investments in network and applications software.

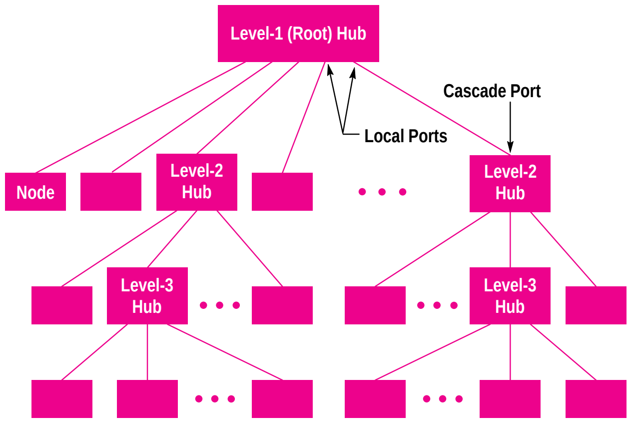

Fig. 1. A cascaded star network.

- Error Susceptibility and Detection. The implementation should have a physical layer bit error rate of less than 10^{-8} and a coding scheme that will guarantee detection of errors in any three bits within the data frame without compromising the cyclic redundancy check (CRC) defined for the IEEE 802.3 and 802.5 frame formats.
- Privacy. Both Ethernet and token ring provide address filtering at the node level, effectively allowing any node to copy any data packet sent on the network. To enhance the privacy of data communications, address filtering of individually addressed packets should be provided as an option within the hub.
- Robust Operation. Continuous operation of the network is required. Physical connections should be tested before allowing nodes to enter the network and provision should be made to allow the identification and removal of disruptive nodes.
- Network Management. An optional network management capability should be provided to monitor network performance, isolate faults, and control network configuration.

Demand Priority and 100VG-AnyLAN

The protocol for the new network is demand priority. It combines the best characteristics of both Ethernet (simple, fast access) and token ring (strong control, collision avoidance, and deterministic delay).

The HP product is 100VG-AnyLAN. It is based on chip technology developed by HP and AT&T that delivers an effective data rate of 100 Mbits/s over several different link configurations: 4-pair Category 3 unshielded twisted-pair (UTP) cable, 2-pair shielded twisted-pair (STP) cable, or 2-pair fiber-optic cable. UTP, STP, and fiber-optic cable can be intermixed on the same network. 100VG-AnyLAN can also operate with the 25-pair bundled cable that is used in many 10Base-T networks, as long as hubs are not connected through bundled cable.

Network Topology

The basic topology used by 100VG-AnyLAN networks is the star configuration. Each hub has two or more local ports and can optionally have one cascade port for connection to a higher-level hub. Nodes can be user stations, bridges to other networks, LAN analyzers, or lower-level hubs.

A network can contain several levels of hubs interconnected in a cascade as shown in Fig. 1. The topmost hub is designated as the level 1 (root) hub. Hubs in lower levels of the cascade are designated by the number of links between them and the root hub. The level number of any particular hub can be determined by the equation:

Hub Level = 1 + (number of link segments away from the root hub)

Hubs at the same cascade level have the same level number.

The maximum number of nodes that can be connected to the network is dependent on the level and frequency of traffic each node generates. The maximum topology diameters and the number of levels of cascading are limited by the allowable delay between the node and the root hub, and can be calculated for any proposed configuration. When there is only one intermediate hub between the root hub and the node, for example, the maximum distance between a node and the root hub is 4 km. Each additional intermediate hub reduces the hub-to-node distance by 1.0 km, resulting in a maximum of four intermediate hubs and a root-hub-to-node distance of 1 km.

Local hub-to-node distances depend on the type of media used for the link: 4-pair UTP links configured with Category 3 or Category 4 cable and links configured with STP cable should not exceed 100 m. Category-5 links can be up to 150 m long. Fiber-optic links can be even longer: 500 m with 850-nm transceivers and 2000 m with 1300-nm transceivers.

Demand Priority Protocol

Control of a demand priority network is centered in the hubs and is based on a request/grant handshake between the hubs and their connected nodes. Access to the network is granted by the hub to requesting nodes in a cyclic round-robin sequence, based on the priority of the request. Within a priority level, selection of the next node to transmit is determined by its sequential location in the network rather than the time of its request. Data is encoded before transmission and is checked for errors at each intermediate hub and the receiving node. Either IEEE 802.3 or IEEE 802.5 frame format can be used.

Architectural Model

The demand priority protocol contains four sublayers corresponding to the two lower layers of the ISO Open Systems Interconnection (OSI) reference model shown in Fig. 2.

The functions of the OSI data link layer are implemented in two sublayers: the LLC and demand priority MAC sublayers. The upper sublayer in a network node is the IEEE 802.2 logical link control (LLC) sublayer. The media access control (MAC) sublayer provides data formatting and control of packet transmission (or reception) in the transmitting (or receiving) node. The MAC also initiates outgoing control requests and acts on received control indications.

Each hub provides control of its connected star portion of the network. The RMAC sublayer provides a superset of the functions of the node's MAC sublayer (except frame formatting). It selects which node will next be granted permission

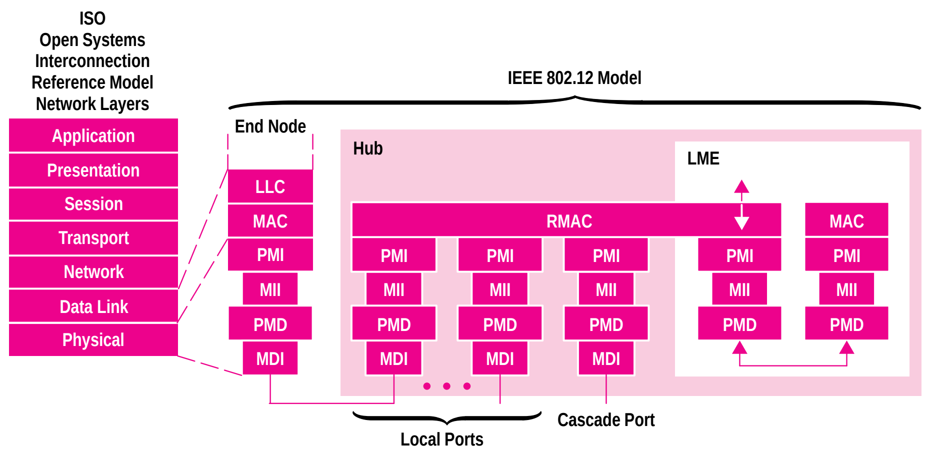

Fig. 2. The demand priority architectural model. LME is the layer management entity.

to send a packet, determines where the received packet will be sent, provides local control of packet reception and retransmission, and monitors each connected link for proper operation. The RMAC initiates outgoing control requests and acts on received control indications. There is no LLC sublayer in a hub.

The functions of the physical layer of the OSI model are also provided by two sublayers: the PMI and PMD sublayers. The physical medium independent (PMI) sublayer presents the same logical interface to both the MAC in a node and the RMAC in a hub. It provides the functions that are common to all link media: ciphering and encoding data and inserting stream headers and trailers before transmission, and removing stream headers and trailers and decoding and deciphering data during reception.

The physical medium dependent (PMD) sublayer contains the functions that are dependent on the particular link medium: bit encoding, signal conditioning, and multiplexing (if necessary) before transmission, and signal recovery, demultiplexing (if necessary), and bit decoding during reception. The PMD translates control requests and generates outgoing control signals. It also detects control signal transitions and generates the appropriate control indication for the MAC or RMAC.

The medium independent interface (MII) is defined in the draft standard as an optional, physically exposed connection to allow link configuration interchangeability (for example, changing from 4-pair UTP to a fiber link). The medium dependent interface (MDI) is the connector between the PMD and the link media.

Standards Development

The demand priority protocol is currently in the process of becoming the IEEE 802.12 standard. The following milestones had been accomplished by May 1995:

- November 1992. A proposal was made for a demand priority development project.
- July 1993. IEEE established the 802.12 Demand Priority Working Group.
- November 1993. An initial draft document (D1) was submitted to the 802.12 Working Group for review.
- July 1994. The IEEE 802.12 draft standard (D4) was submitted to the Working Group for ballot.
- January 1995. The IEEE 801.12 draft standard (D7) was submitted to the LAN/MAN Standards Committee for sponsor ballot.
- March 1995. Sponsor ballot was successfully completed.
- May 1995. The IEEE 802.12 draft standard (D8) was submitted to the Review Committee of the IEEE Standards Board. It is anticipated that the Review Committee and the Standards Board will vote on approving IEEE 802.12 in June.

As the base IEEE l802.12 standard approaches completion, the 802.12 Working Group is investigating future enhancements. At the March 1995 meeting, study groups were formed to investigate higher-speed operation (0.4 to 4 gigabits per second) and a PMD for 2-pair Category 5 UTP. Also under discussion are redundancy and full-duplex links.

100VG-AnyLAN Products

HP currently offers a variety of products for 100VG-AnyLAN within the HP AdvanceStack networking family:

- HP AdvanceStack 100VG Hub15 (HP J2410A) is a 15-port 100VG-AnyLAN hub (Fig. 3).
- HP AdvanceStack 100VG SNMP/bridge module (HP J2414A), when installed in the expansion slot of 100VG Hub15, adds SNMP network management and bridging to 10-Mbit/s Ethernet networks.

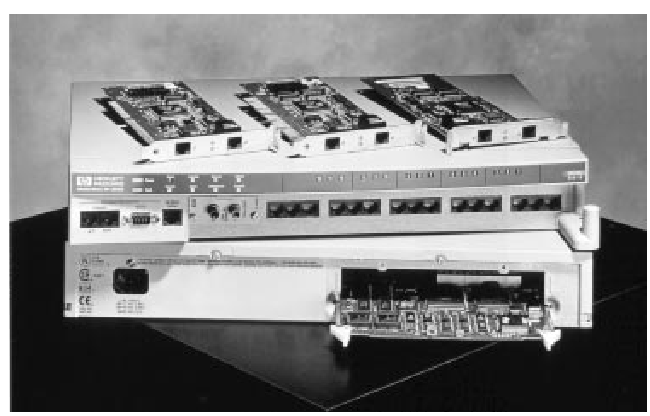

Fig. 3. The HP AdvanceStack 100VG Hub15 (HP J2410A) hub is a 15-port 100VG-AnyLAN hub.

Other Network Technologies

FDDI. Fiber Distributed Data Interface (FDDI) networks run at a standard 100-Mbit/s transfer rate using multimode optical-fiber cabling. The purchase cost is high, partly because it necessitates new cabling for existing networks. FDDI is also available as a high-speed backbone network connecting LANs.

A twisted-pair version, TP-PMD (twisted-pair physical medium dependent) FDDI, is under development, using Category 5 UTP and STP copper cables. This version is sometimes called Copper Distributed Data Interface (CDDI).

ATM. Asynchronous transfer mode (ATM) is a new network technology particularly suitable for wide area networks and campus backbones. It is intended to allow seamless integration of campus LAN backbones into the wide area network.

ATM uses cell switching (53 bytes per cell) similar to high-speed telephone switching over existing UTP or optical-fiber cabling. It runs at 25 to 622 Mbits/s.

Table I
Cabling and Topological Comparisons

	FDDI	ATM	100Base-T	Demand Priority
100-Mbit/s Category 3 cable supported?	No	50 Mbits/s with complex coding	Yes	Yes
Bundled cables supported?	No	No	No	Yes
Multiple cascades supported without bridging or routing?	Yes	Yes	No	Yes
Cost	High	High	Medium	Low

Ethernet and Token Ring Switching. Switching is also being introduced into Ethernet and token ring networks. To increase overall throughput, LANs are segmented. Hubs switch packets dynamically between connected segments allowing simultaneous transmissions among pairs of network segments. This increases bandwidth by two or more times that of individual segments.

100Base-T. 100Base-T is a scaling of CSMA/CD to 100 Mbits/s. There is no migration path or accommodation for existing token ring users. The technique cannot emulate 10Base-T topologies since the maximum topology is two repeaters.

Table II
System Comparisons

	FDDI	ATM	100Base-T	Demand Priority
Supports multimedia with guaranteed delay and bandwidth?	Yes	Yes	No	Yes
End-node adapter card complexity	Node management is expensive	Segmentation or reassembly of frames is expensive	Low	Low
Ethernet 802.3 networks can be upgraded without software changes?	No	No	Yes	Yes
Token ring 802.5 networks can be upgraded without software changes?	No	No	No	Yes

- HP 10/100VG selectable ISA, EISA, and PCI adapters (HP J2573A, J2577A, and J2585A) are PC LAN adapter cards with one RJ-45 connector for 10-Mbit/s 10Base-T and another for 100-Mbit/s 100VG-AnyLAN. The adapters automatically sense which network they are connected to and select the correct mode of operation.
- HP 100VG-AnyLAN/9000 (HP J2645AA, J2655AA) are adapters for HP 9000 Series 700 workstations. They support connection to either 10-Mbit/s 10Base-T or 100-Mbit/s 100VG-AnyLAN. The adapters automatically sense which network they are connected to and select the correct mode of operation.
- HP 100VG-AnyLAN development system (HP E2463A) is a development system for designers and operators of 100VG-AnyLAN network products. It finds the root cause of any IEEE 802.12 design and interoperability problems.

Upgrading Existing Networks

The current 100VG-AnyLAN product set provides a smooth, step-by-step way for customers who wish to upgrade their existing 10Base-T Ethernet networks. The only required elements are new hubs for the network and new adapter cards for each node. Existing network management interfaces, operating systems, bridges, and routers can remain.

In most cases, not all of the network will have to be upgraded at the same time. Consider, for example, an existing 10Base-T network with a mix of high-traffic and low-traffic users that all need to access file servers and printers. The first phase is to identify the high-traffic users to determine how many nodes need to be upgraded. This will determine the number of 100VG-AnyLAN hub ports and network adapters that will be required (including network file servers and printers).

The second phase is to acquire the necessary 100VG-AnyLAN hubs and network adapter cards (one hub for each 15 nodes if HP AdvanceStack 100VG-AnyLAN hubs are used), and to install them in the network as depicted in Fig. 4. The link cable is disconnected from each node that has been identified for upgrading, a replacement PC LAN adapter card is installed, and the link cable is reconnected.

100VG-AnyLAN hubs are installed adjacent to the existing 10Base-T hubs that service high-traffic users. The network cables to each upgraded node are disconnected from the 10Base-T hub and connected to a 100VG-AnyLAN hub (some rearrangement of the 10Base-T node-to-hub connections

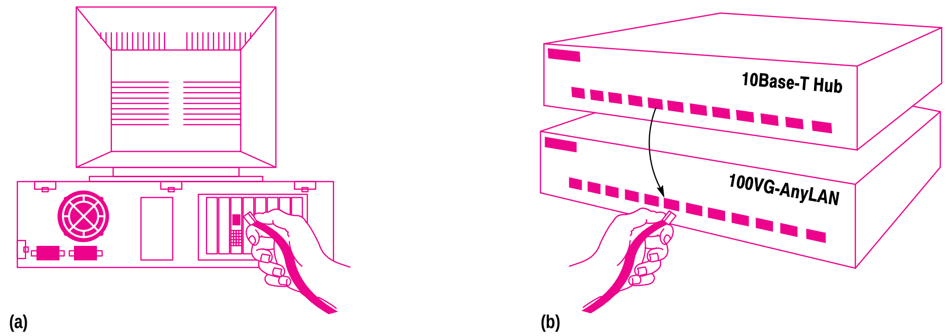

Fig. 4. Upgrading a 10Base-T network to 100VG-AnyLAN. (a) Upgrade adapter at the desktop. (b) Upgrade hub in the wiring closet.

may be advisable). A bridge module between the 100VG-AnyLAN and 10Base-T root hubs interconnects the two LANs.

The network is again ready for use with an upgraded topology as shown in Fig. 5. Additional 100VG-AnyLAN hubs and adapters can be added as needed to accommodate changing traffic levels in the remaining 10Base-T nodes.

The Future

For 100VG-AnyLAN to become a pervasive network topology, the concept of the protocol must be adopted by three major groups: network product manufacturers, network system designers, and network users. Customers must be able to obtain complete sets of the products that are required to meet their individual needs, and the products they obtain should be compatible with each other even if they are obtained from different companies.

An early measure of 100VG-AnyLAN's acceptance can be seen by the number of organizations involved in the demand priority standards development, and in the number adopting the protocol. By January 1995, the following companies had already begun delivering or had announced intention to deliver 100VG-AnyLAN products:

- IC chipsets. AT&T, Motorola, Texas Instruments, Applied Micro Circuits, and Pericom Semiconductor.
- Hubs. HP, Thomas-Conrad, Chipcom, NEC, MultiMedia LANs, Compex, Alfa, Katron Technologies, Optical Data Systems, Anritsu, D-Link, and Ragula Systems.
- Node Adapters. HP, Thomas-Conrad, NEC, Compex, Alfa, Katron Technologies, Ragula Systems, Anritsu, D-Link, Interphase, Optical Data Systems, and Racore Computer Products.
- Multimedia Networking Software. Starlight Networks.
- Development Systems. HP and AT&T.
- Switches and other Internetworking Products. Cisco, Compex, Newbridge Networks, Optical Data Systems, and Plain Tree Systems.
- Desktop Systems. IBM, HP, Compaq, and NEC.

The current 100VG-AnyLAN products provide customers with ten times the speed and up to 16 times the throughput of a 10Base-T Ethernet network at about twice the price. Development of VLSI devices that provide greater integration of the physical layer and the MAC should lead to even lower-cost products.

Demand priority has been designed to be architecturally independent of any particular implementation technology. As such, future generations of demand priority networks may provide higher data rates.

Other Articles in this Issue

The following 100VG-AnyLAN articles provide more detail for their respective areas:

- Demand Priority. The article on page 13 introduces the round-robin pointers, priorities, and bandwidth allocation capabilities of the protocol. A typical demand priority transmission is described in step-by-step fashion, and results of performance simulations are provided.
- Physical Signaling. The article on page 18 gives an expanded description of the physical sublayer. Several design decisions leading to the development of demand priority as a replacement for 10Base-T are described. Quartet signaling, cross talk avoidance, and control signal generation and detection are explained. Differences between the 4-pair UTP and the STP and fiber-optic PMD and link requirements are discussed.

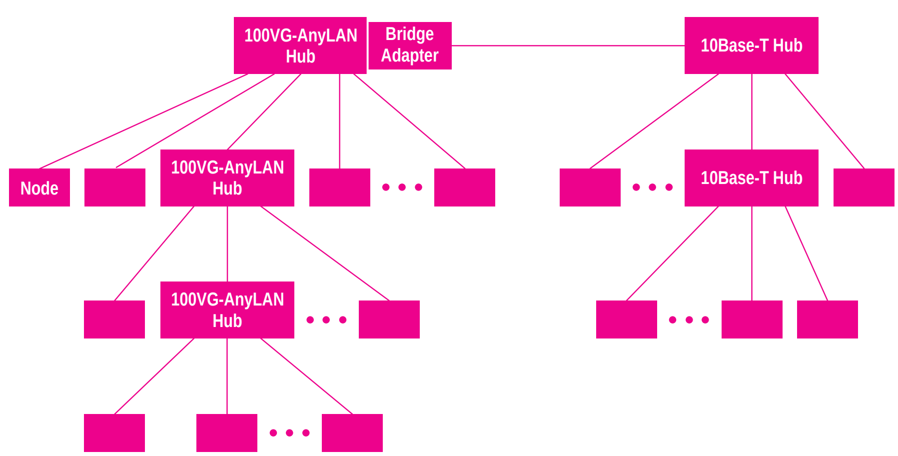

Fig. 5. Hybrid network consisting of interconnected 100VG-AnyLAN and 10Base-T networks.

- Coding. The article on page 27 looks at the techniques of coding data to provide error detection and electrical balance while optimizing the efficiency of the network. Additional insight into the reasons behind the decisions to include both data ciphering and encoding are given along with an expanded discussion of the properties of the coding scheme.
- Multimedia Requirements. The article on page 33 examines the specific demands multimedia applications make on any network system and shows how 100VG-AnyLAN addresses them.

Acknowledgments

The development of 100VG-AnyLAN from public proposal to draft standard and product release has taken only two years. This is an incredibly short time, and has involved almost Herculean efforts of the Roseville and Bristol teams. 100VG-AnyLAN was able to progress from investigation to product in such a short time because of the efforts of a fantastic team. We would particularly like to acknowledge the contributions of our fellow inventors: Alistair Coles, Simon Crouch, David Cunningham, Joe Curcio, Dan Dove, Steve Goody, Jonathan Jedwab, Michael Spratt, and Greg Watson. Although the two halves of the team live 5000 miles apart, these engineers, mathematicians, and physicists collaborated closely to produce a high-performance network. Bill Lane, a professor emeritus from the California State University at Chico, contributed his skills as an editor and writer to produce the text for the draft standard and assisted in the editing of several articles in this issue of the HP Journal.

Completing a major development in a short time requires more than individual commitment. Thus, we also thank our management on both sides of the Atlantic—David Dack, Mark Gasteen, Robert Gudz, Dave Harris, Bill Kind, Gary McAnally, Tim McShane, Carolyn Ticknor, and Steve Wright—for their trust and wholehearted support. They always ensured that we had what we needed to succeed.

References

1. *Information Technology—Local and Metropolitan Area Networks—Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Specifications*, ISO/IEC 8802-3:1993.
2. *Information Technology—Telecommunications and Information Exchange between Systems—Local and Metropolitan Area Networks—LAN/MAN Specific Requirements—Part 5: Token Ring Access Method and Physical Layer Specifications*, ISO/IEC DIS 8802-5:199?.
3. *Local Area Networks—Analysis of the Intelligent Hub Market*, International Data Corporation, May 1994.