The possibility of running Ethernet over UTP excited many companies before the actual standard was developed. Several vendors, including Synoptics,[3] DEC, and others came out with Ethernet-over-UTP variants before the 10Base-T standard was finalized. These variants worked, but naturally, it was impossible to intermix components from these vendors to build an Ethernet-over-UTP configuration. Most of the vendors did, however, have the integrity to guarantee free conversion, or at least plausible migration paths, from their proprietary hardware to hardware compliant with the 10Base-T standard when the latter became available.

This promise does not mean, of course, that all customers bothered to take advantage of such upgrades when they became available. You might find yourself dealing with an existing UTP-based Ethernet installation with which you have had no previous involvement. If so, we suggest you verify that the hardware is indeed 10Base-T compliant rather than an older proprietary variant before blithely assuming that added 10Base-T hardware will integrate properly.

10Base-T Piece Parts 10Base-T NICs are functionally similar to 10Base2 NICs. That is, no separate transceiver is required as with a 10Base5 architecture. The connections are quite different, however. Ethernet UTP wire is terminated at each end with an eight-pin modular RJ-45 jack (only four pins are used). The modular jack plugs directly into the NIC; no T connector is required. Each cable supports only one Ethernet node, as with the pure star thin net approach. See Figure 21-10 for a representation of the 10Base-T environment. At least category 3 wire must be used (see Chapter 43 for more information on categories of UTP wire). Its length must not exceed 100 meters (328 feet).

A 10Base-T hub (sometimes also called a concentrator) is usually required. This is another version of the LAN in a box. The box or hub is the backbone network, whereas the horizontal run to the desktop is the drop. The only exception is if only two devices are to be connected; two 10Base-T devices can be linked via a crossover 10Base-T cable. 10Base-T hubs are functionally identical to multiport repeaters, although some vendors provide sufficiently fast central hubs that the two-repeater rule can be slightly relaxed.

[3] Synoptics later merged with Wellfleet (a router manufacturer), creating Bay Networks. Bay was acquired by Nortel in 1998.

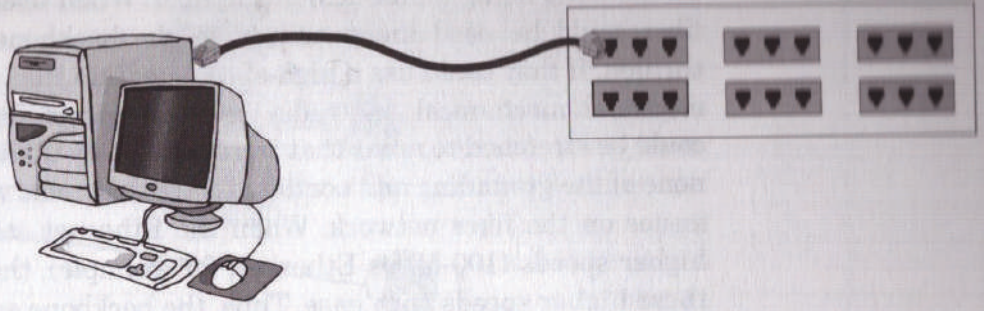


Figure 21-10
The 10Base-T connection

10Base-T hubs (and, to a large extent, other hubs as well) come in two broad categories:

- Work group hubs, sometimes referred to as stackable hubs
- Backplane chassis, sometimes referred to as concentrators

Work group hubs are relatively inexpensive (as low as $200 to $300) devices that typically support 12 or 16 connections. If they incorporate Simple Network Management Protocol (SNMP) management agents, the cost increases by several hundred dollars. Work group hubs, by definition, have very few options in their configurations. Nevertheless, their low cost and simple environmental requirements (they can be mounted in a rack or simply placed on a table or on the floor) make them extremely popular.

Chassis-based concentrators offer great modular flexibility. Such devices consist of a power supply, some control logic, and several slots into which interface boards can be inserted. Each interface board can support several 10Base-T cables, often 12. Depending on the manufacturer, other boards might be available to support different types of networks out of the same chassis. Those other network types (for example, token ring) might be interconnected to the Ethernet board either via an additional interface board just for that purpose or via an external device not inserted into the chassis.

10Base-F (Fiber)

Newer uses of various cabling systems have always been explored. The fiber-optic medium has caught everyone's attention because of the declining costs of the glass and the electrical characteristics (that is, the freedom from electrical and radio frequency interference and the isolation from electrical

components being nonissues) of the fiber. When users discovered that the fiber could be used inexpensively in the backbone network, they were thrilled. If they could use a high-speed medium that was impervious to the electrical, mechanical, and radio frequency interference, then the network could be stretched to areas that were previously unavailable. Furthermore, none of the grounding and bonding issues that come with copper cable were issues on the fiber network. When the Ethernet standard ramped up to higher speeds (100-Mbps Ethernet, for example), the fiber would support these higher speeds with ease. Thus, the backbone saw more implementations with the 10Base-F installed.

The 10-Mbps standard was still used in the backbone network, but UTP wiring was still the least expensive proposition for wiring to the desktop. Therefore, a medium changer is required. In the closets where the LAN will run from floor to floor or from closet to closet, fiber is used. At the hub, ports exist for the connection from hub to hub via the fiber backbone. Inside the hub, the electronics are present to convert the fiber backbone to a copper station drop, or, in more precise terms, from the light to the electrical pulses needed for the copper. Figure 21-11 shows a typical configuration of the 10Base-F environment. This could be extended to the desktop, but, as already stated, the cost of putting fiber cards acting as transceivers in a PC or workstation is too steep for the average user. Therefore, this is done in a mixed environment.

10Broad36

Some networks have been built based on a broadband coaxial cable in the backbone. This might be more prevalent in a campus area (such as a college, corporate park, or hospital) where multiple buildings are interlinked. The use of the broadband cable was an expedient to support the needs of the organization's voice, data, and video needs because all of these components were analog transmissions that were connected between buildings. Therefore, the broadband backbone cable (such as a CATV cable) was used when the Ethernet was introduced. Customers who needed to link multiple buildings together were constrained by the 500-meter distance limitations of the Ethernet baseband coax. Even with the two-repeater rule, the most that the cable could be extended was 1,500 meters, just under a mile. Although this might sound like a lot of cable, which should reach just about anywhere, the one mile of cable gets used very quickly when it needs to be snaked through ceilings, up and down from the ceiling heights, and so on.

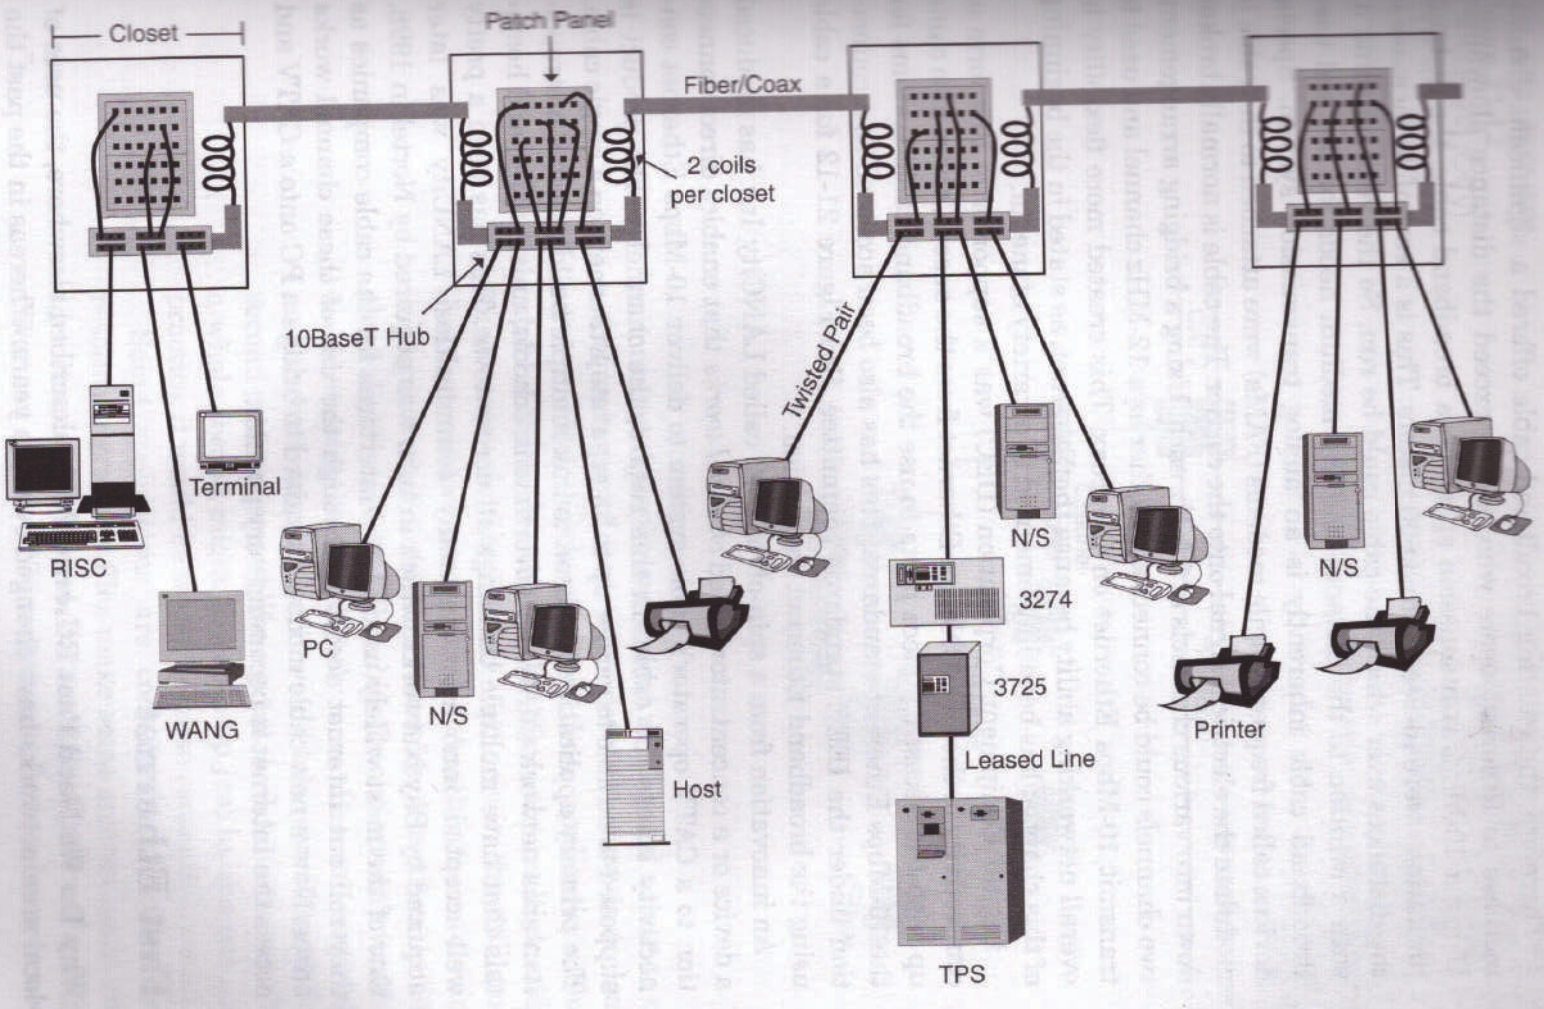


Figure 21-11
The 10Base-F uses fiber in the backbone, but converts to copper to the desk.

Therefore, the existing broadband cable offered a significant distance increase to Ethernet users who had to exceed the distance allowances. Using a 10-Mbps transmission speed on a broadband cable, the distance limitations were increased to 3,600 meters. This is a sevenfold increase in the distances over which the cable could be run. So the vendors came up with a solution to the distance and the medium needs for the end user. Broadband cable inherently is an analog transmission system. Special devices called frequency agile modems (FAMs) were attached to the cable to modulate the electrical signal onto the carrier. The cable is normally broken down into various channels of 6 MHz each. Using a bridging arrangement, two channels could be connected together as a 12-MHz channel and used to transmit 10-Mbps Ethernet on this coax. This created more flexibility in overall networking ability because the Ethernet, as stated in the beginning of this chapter, has been implemented in a variety of media.

Digital Equipment Corporation (DEC) was a supporter of this form of connectivity in their DECNet on Ethernet from the onset. Others also came up with the necessary piece parts to use the broadband cable systems for the 10-Mbps Ethernet standards. This has also been accepted as a connection under the IEEE standards committee. See Figure 21-12 for a cable using the broadband Ethernet connection.

An innovation from a spin-off of DEC called LANCity, Inc. has produced a device or a concentrator called *channel works* that enables direct connection to a CATV operator's cable system to deliver 10-Mbps Ethernet connectivity across the cable. In this regard, the channel works product is supposed to be able to garner up to 83 separate Ethernets on a single cable. The primary application has been to link multiple buildings in a metropolitan area network (MAN) for government, schools, universities, and hospitals that have multiple buildings all across town. This has been a pretty well-accepted innovation in many communities. LANCity was later acquired by Bay Networks, which in turn was acquired by Nortel in 1998. One of the most widely used CATV interfaces for the cable companies as they roll out Internet access is through the use of these channel works boxes. Now a new cable modem standard to bridge a PC onto a CATV and access the Internet is becoming very popular.

Fast Ethernet

Why Do We Need Fast Ethernet? As described elsewhere, the uses of local area networks have changed over the years. Whereas in the past the

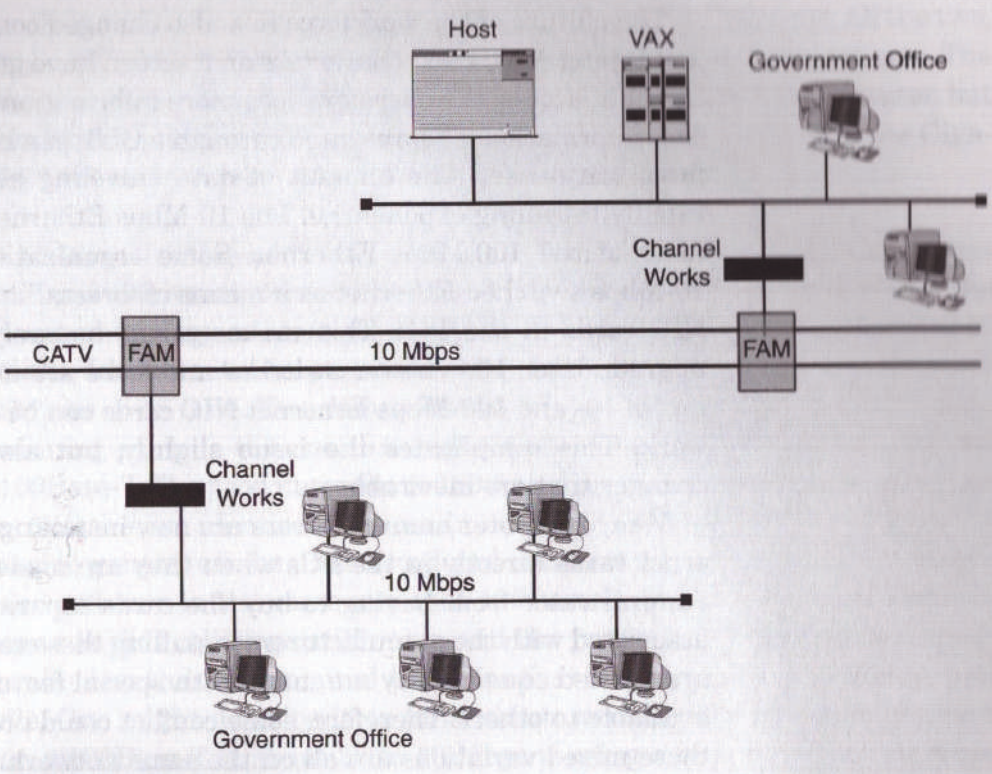


Figure 21-12
A broadband (CATV) application using 10 Mbits/s on coax

primary use was for movement of files and access to data by applications, a frequent use of local area networks now is to load applications from servers. Also, the character of those applications has changed. The executables themselves have moved into the multiple megabyte range, and in some cases, graphic files are also now in the millions of bits. This means that faster networks are needed. Another factor driving the need for speed is larger, higher-resolution monitors. A modern 1,024 × 768 monitor displays slightly over 1 million bits for one screen's worth of information. That would require, ideally, slightly more than one-tenth of a second to traverse an Ethernet. In practice, it would take considerably longer. One-tenth of a second may not sound like a great deal, but on a small network of 10 users, a whole second might be eaten up just by the one screen to each user. In practice, it would be worse because contention would slow things down.

Some applications are considerably worse than this. Large, high-resolution graphics files make some environments entirely unsupportable by a standard Ethernet. For example, the size of a single digitized magazine cover can approach 70 megabytes. Bottom line: we need more speed.

The culture of the workforce has also changed considerably. The days of just typing ASCII text characters on a screen have gone. With Internet and Intranet access, the user now does more information processing. Moreover, that information is represented through a GUI, in which the graphics interfaces are denser. The amount of data traveling across our networks is rapidly becoming exponential. The 10-Mbps Ethernet was pushed aside in favor of the 100-Mbps Ethernet. Some organizations have moved to a 10-Mbps switched Ethernet as a means of forestalling the change of all the NIC cards in the PCs. This makes sense, but only begs for the actual upgrade later. Now newer switches and hubs are auto-detecting where a mix of 10- and 100-Mbps Ethernet NIC cards can be used on the same network. This complicates the issue slightly, but also gains time for the changes that are inevitable.

Many computer manufacturers are now installing 10- to 100-Mbps Ethernet cards directly in the PCs when they are made. This saves the LAN administrator from having to buy the cards separately. A minor issue is associated with the manufacturers installing these cards. Not all NIC cards are created equal. Many are made with special features and functions not available to others. Therefore, some conflict could occur on the LAN when these mixed variations are all on the same network. The other problem is that the computer manufacturers who install the NICs usually require special drivers for their machine. Thus, a user must be aware that differences will arise and the price we pay for free LAN cards is the added burden of keeping on top of the myriad protocols and drivers required.

As the convergence of voice and data caught on in the industry during the late 1990s, the move is to incorporate voice, data, and streaming video at the desktop on the LAN. Given the limits of the 10-Mbps Ethernet with the specifications discussed earlier, the movement to 100 Mbps was escalated.

The term Fast Ethernet is primarily used today to refer to a technology described as 100Base-T. In 1996, when the technology was introduced, Fast Ethernet actually referred to several technologies. This creates some confusion. One of the technologies was switched Ethernet. Switched Ethernet is described in another chapter, but generally speaking, Fast Ethernet no longer refers to switched 10-Mbps Ethernet. Likewise, full-duplex Ethernet, at least the full-duplex variation of 10Base-T, is no longer referred to as Fast Ethernet. Just as with 1,200-bps modems, which once were described as high-speed, the advent of better technologies has resulted in changing the terminology used to describe the older technologies.

Only two technologies really deserve the name Fast Ethernet. Of the two, the vast majority of analysts have declared 100Base-T the winner. The other is 100VGAnyLAN. We'll discuss both technologies in this chapter, but make no mistake; 100VGAnyLAN is dead. We'll also briefly describe Gigabit Ethernet, a technology that was standardized in 1998.

100Base-T When developers were trying to determine what the successor to Ethernet should be, they had two basic choices: try to stick to the basic technology as closely as possible, just increasing its speed, or try to make fundamental improvements to solve flaws that had become more obvious over time. Two camps developed, one taking each approach. The winning camp is the one that took the first approach. Their product was 100Base-T. The good news about 100Base-T is that the frame size and characteristics are identical to those in standard Ethernet. It uses CSMA/CD and in general can be understood and managed by network administrators and software in the same way that 100Base-T could be managed. The bad news is that collisions still occur, and the product still taps out at considerably less than the nominal 100 Mbps that it is supposed to be able to handle. One might describe it as *Ethernet on steroids.* The actual throughput on a 100Base-T is about 33 to 40 Mbps if using a hub and about 60 Mbps if using an Ethernet switch.

It is, however, the most popular replacement for standard Ethernet by far. This is due largely to good marketing, but something should definitely be said for not having to do major retraining of your network administrators and technicians. In 1998, 50 percent of all new LANs installed were at the 100-Mbps rate. In 1999, that number jumped to almost 67 percent. This varied between the shared bus concept in a hub and the switched version using a switch. Regardless, the costs, although higher than 10-Mbps NICs, are sufficiently attractive to merit the installation regardless of whether the speed is needed at the desktop today.

Like the original Ethernet, 100Base-T supports several media. Notably absent from this list is coaxial cable. Both standard copper wiring and fiber are supported. As with all high-speed technologies, the copper implementation has severe distance limitations. Most of the high-speed technologies do not support end devices more than 100 meters away from their central electronics locations; Fast Ethernet is no exception.

Fast Ethernet can be handled on three variations of copper wire. The first is category 3 UTP (unshielded twisted pair). If you are going to use category 3, however, you must use four pairs of twisted-pair wiring. This is an

extremely unpopular option. It is not difficult to see why. Although many installations have some copper wiring available, it is extremely unusual to have four pairs that are unused. If one is going to have to put in new wiring, there is no reason to put in anything less capable than category 5. In addition, Fast Ethernet on category 3 only supports half-duplex operation. If Fast Ethernet is used in half-duplex mode on four pairs of either category 3 or category 5 wiring, it is described as 100Base-T4.

The second copper variation supported is category 5 UTP. In the case of category 5, Fast Ethernet can operate in either half or full duplex on two pairs of wire rather than the four required for category 3. Operating on two pairs, Fast Ethernet is referred to as 100Base-TX. This is by far the most popular implementation of Fast Ethernet.

The third variation supported on copper for Fast Ethernet is shielded twisted pair (STP). Generally speaking, STP should only be installed in special cases where the potential interference either to or from the network is an issue. STP is much more difficult to install, with more demanding grounding requirements than UTP. It doesn't increase speed, nor does it improve usage of the pairs in the cable. As far as network performance is concerned, it works the same as on UTP.

In addition to copper wiring, 100Base-T is supported on fiber. As with category 5, fiber supports both half- and full-duplex operations. Why might you use fiber? It doesn't improve the speed. The reason is fiber's general superiority over copper, as described elsewhere in this book. In the case of Fast Ethernet, the main reason to use fiber would be to extend the available distances. Fast Ethernet over copper, as mentioned earlier, only reaches 100 meters to the desktop. A more severe limitation is that distances between copper-connected Fast Ethernet hubs must not exceed 5 meters (Figure 21-13). The use of fiber considerably relaxes these limitations.

As with standard Ethernet, the way the wiring (either copper or fiber) can be used varies. Buyers must select among shared, switched, and/or full-duplex use of the media and obtain the right electronics to install at both ends of the cable. The default is shared half duplex. One should be aware that just as in standard Ethernet, the contention experienced with a shared half-duplex configuration means that in no case, except a two-user network, will you be able to count on actually getting anywhere near the 100-Mbps nominal performance of a Fast Ethernet. However, that should not usually be a major problem. Just as with standard Ethernet, medium or even large work groups can be set up that still get a significant performance improvement over standard Ethernet by using shared half-duplex Fast Ethernet.

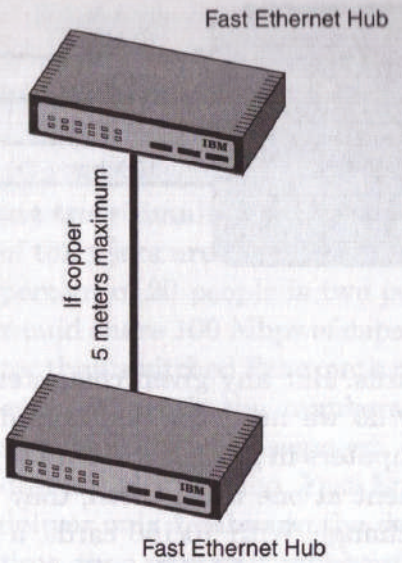


Figure 21-13
Copper-connected Fast Ethernet hubs have to be close!

If what you're doing is building a backbone, you probably want to use switched full-duplex Fast Ethernet. This will deliver close to the rated 100-Mbit performance to each node, and it costs considerably less than the available alternatives (FDDI and ATM are the ones that come to mind). Note that for anything but shared half duplex, one must use either category 5 UTP or fiber cabling. Naturally, the cost of the electronics involved to deliver switched or full duplex is considerably more than that for shared half duplex—in fact, three to five times the cost of the less expensive approach.

Changing your network over to Fast Ethernet has been made much easier than it might have been by clever manufacturers that have addressed the problem. Because all Fast Ethernet implementations are based on either copper or fiber, the cabling is done in a star or radial configuration; no physical bus topologies need apply. This means that in theory and in practice, each node can be upgraded individually. The manufacturers build the central electronics (hubs, switches, and so on) with the capability, in most cases, to handle either 10- or 100-Mbps Ethernet on a port-by-port basis. This means that better electronics can be swapped at the central location, the network can be brought back up, and one-by-one, individual users' machines can be upgraded (Figure 21-14).

Better than this, however, is the advent of dual-speed NICs. Several manufacturers make NICs that operate at both 10 and 100 Mbps. These are

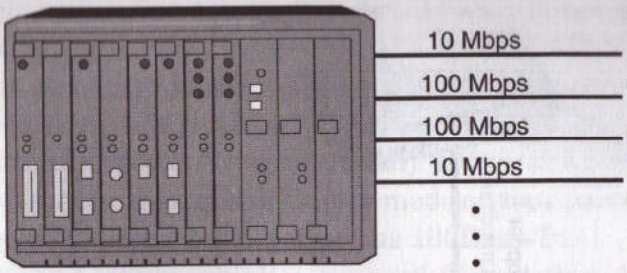


Figure 21-14
Implementing Fast Ethernet — 10/100 hub

described as 10/100 cards. But any given computer can run at only one speed at a time—why do we need the dual speed? Most environments upgrade individual computers in groups; it is rare to replace all of the computers in an environment at one time. Often, they are replaced one at a time, as users' needs change. With 10/100 cards, a company can make a long-term decision to move toward 100Base-T without having to upgrade all the central hubs. The company can upgrade the desktops and then upgrade the hubs as it makes sense to do so. In most cases, the 10/100 cards automatically sense the speed of the connected wire (that is, of the central electronics port) at the time the PC is booted.

A changeover plan could be developed in one of two ways. One plan would be to replace all of the NICs in desktops over time with 10/100 cards. These cards are not significantly more expensive than standard Ethernet cards. Once all of the cards are replaced, then replace the central electronics with 100Base-T switches and hubs, reboot all the desktops, and you are done. The second approach would be to replace the central electronics with switches and hubs with units that have 10/100 ports. These are more costly, but give a degree of flexibility that isn't available in the first scenario; not all desktops have to be upgraded either. Once the central electronics are all upgraded, one can begin upgrading the desktops. In a pinch, some desktops can be upgraded in parallel, although this complicates the management efforts.

100Base-T Performance Expectations The most obvious alternative to implementing 100Base-T is switched 10Base-T Ethernet. Surprisingly, the latter may, under some circumstances, deliver better performance than shared 100Base-T. One may think that this is an unfair comparison; however, the costs are similar, whereas the cost of switched 100Base-T is probably much too high for desktop use in most cases. So we will consider switched 10Base-T and shared 100Base-T as alternatives.

It all boils down to the numbers. In an environment with small workgroups (fewer than 20), shared 100Base-T will probably deliver much better performance to the desktop than shared Ethernet could. Why? With 100Base-T, each user will average at most 5 Mbits of capacity on the hub if all are active simultaneously. But in most real-world scenarios, not all users are truly simultaneously active. In the typical case, fewer than 10 percent of the users are doing anything on the network at any given instant. Ten percent of 20 people is two people. These two active users, in most cases, would share 100 Mbps of capacity, or 50 Mbps each, which is obviously better than switched Ethernet's delivery of 10 Mbps at best to each user. However, let's cook the numbers a little differently. If 100 users are on a switched Ethernet segment, each will still get 10 Mbps delivered to the desktop all the time. Fast Ethernet at 100 Mbps, in the worst case, will deliver only 1 Mbps to the desktop. If 10 percent of users are active at a time, then each will get about 10 Mbps. In practice, contention on the 100-Mbps hub will reduce that by 20 to 30 percent. In this case, switched Ethernet might be a better choice than the Fast Ethernet.

A major decision factor is the type of work the users are doing. If each is doing network-intensive activities on a continuing basis that can be satisfied by a 10-Mbps pipe, switched Ethernet is probably the technology of choice. On the other hand, if the user population makes sporadic use of the network, but each use tends to benefit from a truly high bandwidth (editing large graphic files is an application that comes to mind), then 100 Mbps (or higher) would definitely be the better choice because each user would experience close to the maximum capacity when needed, if it's only lightly used.

100VGAnyLAN Back before 100Base-T had established its market share, a group of technically oriented companies including Hewlett-Packard felt that the second approach to upgrading Ethernet—improving its flaws as opposed to just increasing its speed—was the better approach. The result of their efforts was 100VGAnyLAN. This technology, like the Betamax VCR, is clearly superior to its competitor (100Base-T). Unfortunately, as with the Betamax, it didn't win over the market. One can still buy 100VGAnyLAN products; however, the system is regarded by most as an also-ran.

100VGAnyLAN had much to recommend it. Its designers chose to eliminate the CSMA/CD method of allocating capacity. Instead, it uses what is called a demand-priority allocation scheme, where each NIC notifies its corresponding hub that it has something to send. The hub pays attention to these notifications and solicits the input once notified. It is also possible to

designate high-priority traffic, such as video, enabling a better pattern of performance for such traffic in this technology.

100VGAnyLAN also had the capability to support either Ethernet or token ring traffic. This traffic cannot be on the same hub, but one can use the same kind of hubs for both, significantly reducing the complexity of a LAN implementation environment. The rules for cascading 100VG hubs are no more restrictive than those for standard Ethernet and considerably more flexible than those for 100Base-T. As with 100Base-T, the cable running out to the workstations can be 100 meters long; however, the distance between hubs also can be 100 meters, 20 times the limitation of 100Base-T.

100VGAnyLAN has some disadvantages, partially explaining its failure in the marketplace. For one thing, its initial specification required four pairs of unshielded twisted pair wiring, no matter which category was being used. Another disadvantage is more political than technological. Since the beginning, Ethernet has used CSMA/CD as its media allocation scheme, 100VGAnyLAN does not. It uses demand priority. In fact, this is an advantage; however, purists in the Ethernet camp have attacked 100VGAnyLAN on this basis. It is a poor reason not to use the technology; however, enough advocates of Ethernet went the other way because of this that it seriously and adversely affected 100VGAnyLAN's market share. Very few vendors subscribed to the 100VGAnyLAN technological specification. At this point, it is really only enthusiastically supported by Hewlett-Packard, and even HP has moved on to support 100BaseT.

Gigabit Ethernet Although we do not know of anyone who alleges that full-duplex switched 100Base-T is insufficient to handle any likely desktop requirement, if one puts together enough high-speed desktop LANs, the backbone to serve them needs to be a better technology than that desktop LAN technology, whatever it is. Moreover, if the use of voice, data, video, and multimedia finally becomes a reality on a single architecture, then the need for more than 100 Mbps to the desk may become a reality. Gigabit Ethernet was developed with the goal of providing these applications to the individual user in the future. In 1998, products were introduced before the standards had been adopted. As with many new technologies, some vendors came out with products early and promised to change or upgrade those products, at no cost to the initial implementers, to whatever the standard becomes. This allowed the vendors to get a leap on the market and assure their market share before others could get there. So the question is, when would you use Gigabit Ethernet and when would you not?

Like 100Base-T, Gigabit Ethernet comes in half- and full-duplex variations. In its half-duplex implementation, it has some very severe limitations regarding delivery of actual capacity. An interaction occurs in all Ethernets between minimum packet size and size in terms of radius of the actual maximum network configuration. The smaller the minimum frame size, the smaller the maximum radius of the network. This relationship is compounded in a bad way by increasing the speed of the network transmission. Because the goal in upgrading to a faster variety of Ethernet is to keep the same frame size limitations, the radius of the Gigabit Ethernet without using a special technique is limited to about 20 meters, which is obviously impractical. The designers implemented a special technique that involves transmission of a special signal in any case where a frame smaller than 512 bytes is transmitted on a Gigabit Ethernet. This means that for large frames (those associated with large file transfers), Gigabit Ethernet might actually deliver close to its full speed. However, many networks being used by human beings as opposed to machines generate a large number of small frames averaging only 300 to 350 bytes. In the worst case, such a network, even if implemented on Gigabit Ethernet, would experience performance a little better than that delivered by 100Base-T in its full-duplex incarnation. The technique creates a 200-meter network diameter. This limitation means that any network implementer needs to seriously analyze what kind of traffic is going to be carried by that network and evaluate the likely bandwidth to actually be delivered by Gigabit Ethernet.

Gigabit Ethernet also did not provide quality of service (QoS) guarantees, although the proponents of Gigabit Ethernet are working to provide QoS-like services. This means that even though its nominal speed is very high, it still is not an ideal technology for satisfying isochronous type traffic (such as video). The new standards for the 802.3Q are designed to handle QoS features for Gigabit Ethernet.

If one implements full-duplex Gigabit Ethernet, the issue regarding distance pretty much disappears. Gigabit Ethernet requires the use of fiber. Two kinds of fiber are available: multimode and single mode. Multimode should be able to reach up to about 550 meters, whereas single mode should be able to reach distances approaching 3 kilometers.

One other caution: Most of the Intel-based servers in service are generally incapable of accepting or transmitting information at speeds significantly greater than those that can be provided by 100Base-T. Thus, hooking a server directly to a Gigabit pipe would not necessarily deliver improved performance over Fast Ethernet. Higher-end UNIX servers can deliver

performance two to four times the above, but even they cannot yet receive or transmit information at true Gigabit speeds. Backbone components (routers, and so on) also are just now coming on the market that can handle speeds in this range. Therefore, simply buying Gigabit devices and dropping them into the network is not likely to be a successful strategy.

Gigabit Ethernet does promise to be a less expensive technology than ATM. In cases where backbones can be successfully implemented with Gigabit Ethernet, it may be more cost-effective. However, ATM is available at speed ranges that are at least competitive with Gigabit Ethernet. With the additional capabilities of ATM and its large (carrier-based) installed base, ATM is probably the preferred technology for high-capacity backbones at least for the next couple of years.

Just when we thought it was safe to come out of the LAN closets and use the technologies, with the knowledge that Gigabit Ethernet is coming along to deal with bandwidth issues, the next in a series of evolutionary steps emerged. Now the talk is focused on 10-Gigabit Ethernet. Nortel Networks announced their strategy in 1999 with the use of 10-Gbps Ethernet over Synchronous Optical network (SONET) OC-192-based networks. The 10-Gigabit Ethernet standards are still emerging, so many vendors are coming out with their own proprietary solutions to satisfy this demand. Once again, this places the buyer at risk when the standards finally get approved sometime in 2003-2004.

This 10-Gigabit strategy also points out that the end user may not need this speed yet, but the carriers (telcos, ISPs, and new providers) will be looking for this throughput. If a carrier chooses to offer 1 Gbps to a consumer (large organization), then the 10 Gbps is necessary for the provider to guarantee the QoS and the effective data rates. This is particularly true when serving the ISPs and Internet Telephony Service Providers (ITSPs). Through the use of Metro Area Networks running on SONET standards and using Dense Wave-Division Multiplexing (DWDM) on the fiber systems, the carriers will be delivering 1-2-54 or 10 Gbps Ethernet.

Virtual LANs

In a traditional shared media network, traffic generated by a station is propagated to all other stations on the local segment. For a given station on shared Ethernet, the local segment is the collision domain because traffic on the segment has the potential to cause an Ethernet collision. The local

segment is also the broadcast domain because any broadcast is sent to all stations on the local segment.

Ethernet bridges and switches divide a network into smaller collision domains, but they do not affect the broadcast domain. In simple terms, a *virtual local area network* (VLAN) can be thought of as a mechanism to fine-tune broadcast domains.

A VLAN is a collection of switch ports that make up a single broadcast domain. A VLAN can be defined for a single switch, or it can span multiple switches. VLANs are logical entities created in the software configuration to control traffic flow. Membership in workgroup segments can be determined logically instead of by user location, and moves, adds, and changes can be easily configured as the network evolves.

IEEE 802.11 Wireless LANs

The IEEE 802.11 standard defines the physical layer and media access control (MAC) layer for a wireless local area network (WLAN). The standard defines three different physical layers for the 802.11 WLAN, each operating in a different frequency range and at rates of 1 Mbps and 2 Mbps. In this section, we focus on the architecture of 802.11 LANs and their media access protocols. We'll see that although it belongs to the same standard family as Ethernet, it has a significantly different architecture and media access protocol. The IEEE 802.11 represents the first standard for WLAN products from an internationally recognized, independent organization. The IEEE manages most of the standards for wired LANs. We saw the standards at work when we reviewed the Ethernet (802.3) and the Token Ring (802.5). The WLAN standard (802.11) represents an important milestone in WLAN systems because customers can now have multiple sources for the components of their WLAN systems. Some applications are still available where the existing proprietary data communications are a good fit because they may optimize some aspect of the network performance. However, 802.11 compliant products expand the users' options.

802.11 LAN Architecture The principal components of the 802.11 WLAN architecture are shown in any architectural drawing. The fundamental building block of the 802.11 architecture is the cell, known as the *basic service set* (BSS) in 802.11 parlance. A BSS typically contains one or more wireless stations and a central base station, known as an *access point*

(AP) in 802.11 terminology. The stations, either fixed or mobile, and the central base station communicate among themselves using the IEEE 802.11 wireless MAC protocol. Multiple access points may be connected together (using a wired Ethernet or another wireless channel) to form a so-called *distribution system* (DS). The DS appears to upper-level protocols (such as, IP) as a single 802 network, in much the same way that a bridged, wired 802.3 Ethernet network appears as a single 802 network to the upper-layer protocols.

How It Will Be Used in End Applications? The IEEE 802.11 standard defines the protocol for two types of networks: ad hoc and client/server networks.

- An *ad hoc network* is a simple network where communications are established between multiple stations in a given coverage area without the use of an access point or server. The standard specifies the etiquette that each station must observe so that they all have fair access to the wireless media. It provides methods for arbitrating requests to use the media to ensure that throughput is maximized for all of the users in the base service set. IEEE 802.11 stations can also group themselves together to form an ad hoc network—a network with no central control and with no connections to the outside world. In this case, the network is formed on the fly simply because mobile devices have found themselves in proximity to each other, have a need to communicate, and find no pre-existing network infrastructure (such as a pre-existing 802.11 BSS with an AP) in the location. An ad hoc network might be formed, for example, when people with laptops meet together (in a conference room, a train, a car, and so on) and want to exchange data in the absence of a centralized access point. A tremendous recent increase in interest in ad hoc networking has occurred, as communicating portable devices continue to proliferate. Figure 21-15 illustrates the ad hoc network.
- The *client/server network* uses an access point that controls the allocation of transmit time for all stations and enables mobile stations to roam from cell to cell. The access point is used to handle traffic from the mobile radio to the wired or wireless backbone of the client/server network. This arrangement enables point coordination of all of the stations in the basic service area and ensures proper handling of the data traffic. The access point routes data between the stations and other wireless stations or to and from the network server. Typically,

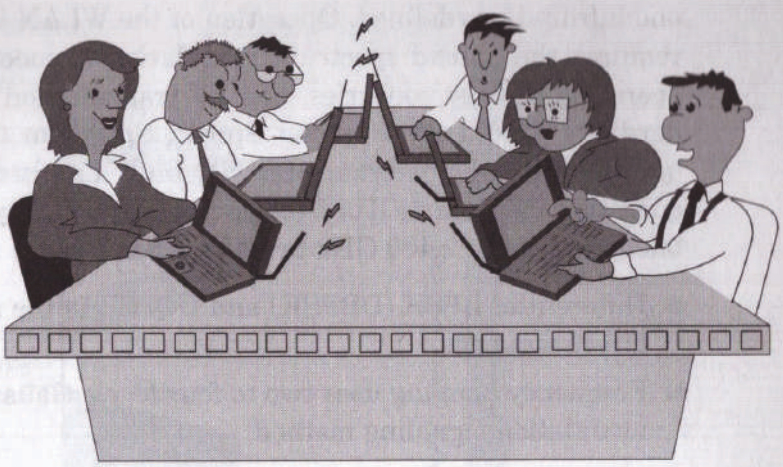


Figure 21-15
The ad hoc network

WLANs controlled by a central access point will provide better throughput performance. Figure 21-16 illustrates the client/server approach.

Physical Layer Implementation Choices The physical layer in any network defines the modulation and signaling characteristics for the transmission of data. At the physical layer, two RF transmission methods and

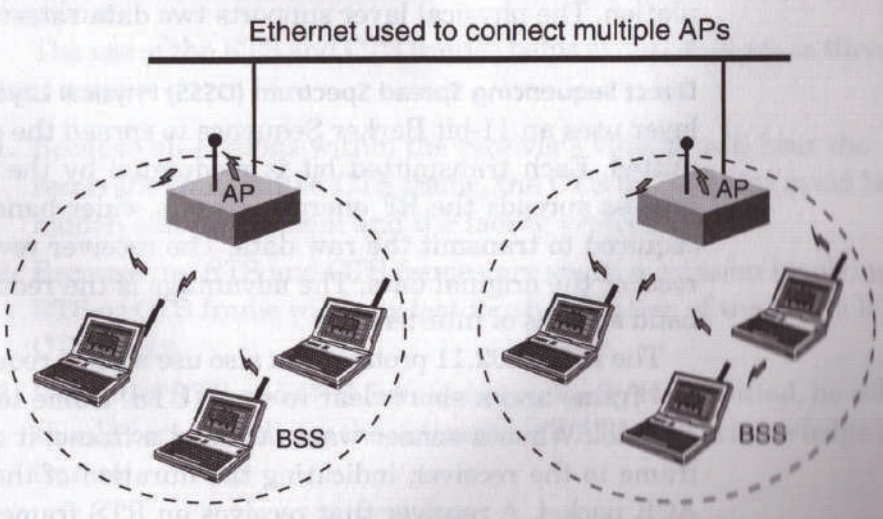


Figure 21-16
The client/server approach to wireless LANs

Legend:
AP= Access Point
BSS= Basic Service Set

one infrared are defined. Operation of the WLAN in unlicensed RF bands requires the spread spectrum modulation to meet the requirements for operation in most countries. The RF transmission standards in the standard are Frequency Hopping Spread Spectrum (FHSS) and Direct Sequence Spread Spectrum (DSSS). Both architectures are defined for operation in the 2.4-GHz frequency band typically occupying 83 MHz of bandwidth from 2.400 GHz to 2.483 GHz.

- Differential BPSK (DBPSK) and DQPSK is the modulation for the direct sequence.
- Frequency hopping uses two to four levels Gaussian FSK as the modulation signaling method.

Different frequencies were approved for use in Japan, the United States, and Europe and any WLAN product must meet the requirements for the country where it is sold. The physical layer data rate for FHSS system is 1 Mbps. For DSSS, both 1-Mbps and 2-Mbps data rates are supported. The choice between FHSS and DSSS will depend on a number of factors related to the user's application and the environment that the system will be operating.

Infrared Physical Layer One infrared standard is supported which operates in the 850nM- to 950nM-band with peak power of 2 W. The modulation for infrared is accomplished using either 4- or 16-level, pulse-positioning modulation. The physical layer supports two data rates: 1 and 2 Mbps.

Direct Sequencing Spread Spectrum (DSSS) Physical Layer The DSSS physical layer uses an 11-bit Barker Sequence to spread the data before it is transmitted. Each transmitted bit is modulated by the 11-bit sequence. This process spreads the RF energy across a wider bandwidth than would be required to transmit the raw data. The receiver reverses the RF input to recover the original data. The advantage is the reduced effects of narrowband sources of interference.

The IEEE 802.11 protocol can also use a short request to send (RTS) control frame and a short clear to send (CTS) frame to reserve access to the channel. When a sender wants to send a frame, it can first send an RTS frame to the receiver, indicating the duration of the data packet and the ACK packet. A receiver that receives an RTS frame responds with a CTS frame, giving the sender explicit permission to send. All other stations hearing the RTS or CTS then know about the pending data transmission and can avoid interfering with those transmissions. An IEEE 802.11 sender can

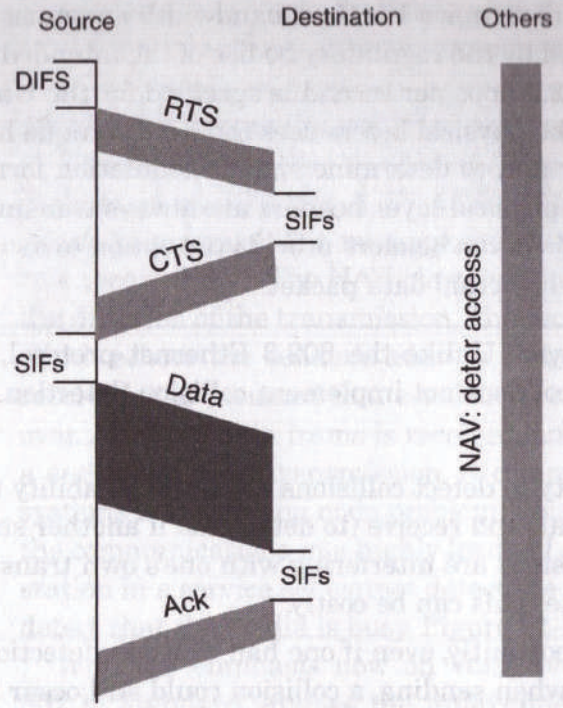


Figure 21-17 The RTS/CTS sequence

operate either using the RTS/CTS control frames or can send its data without first using the RTS control frame. Figure 21-17 shows the RTS/CTS sequence.

The use of the RTS and CTS frames helps avoid collisions in three important ways:

1. Because all stations within the receiver's vicinity will hear the receiver's transmitted CTS frame, the CTS frame helps avoid both the hidden station problem and the fading problem.
2. Because the RTS and CTS frames are short, a collision involving an RTS or CTS frame will only last for the duration of the whole RTS or CTS frame.
3. When the RTS and CTS frames are correctly transmitted, no collisions should occur involving the subsequent DATA and acknowledge (ACK) frames.

Frequency Hopping Spread Spectrum (FHSS) Physical Layer The FHSS physical layer has 22 hop patterns to choose from. The frequency hop physical layer is required to hop across the 2.4-GHz ISM band covering 79 channels.

Each channel occupies 1 MHz of bandwidth and must hop at the minimum rate specified by the regulatory bodies of the intended country. A minimum hop rate of 2.5 hops per second is specified for the United States.

Each of the physical layers uses their own unique header to synchronize the receiver and to determine signal modulation format and data packet length. The physical layer headers are always transmitted at 1 Mbps. Predefined fields in the headers provide the option to increase the data rate to 2 Mbps for the actual data packet.

The MAC Layer Unlike the 802.3 Ethernet protocol, the wireless 802.11 MAC protocol does not implement collision detection. This is because of a few reasons:

- The ability to detect collisions requires the ability to both send (one's own signal) and receive (to determine if another station's transmissions are interfering with one's own transmission) at the same time. This can be costly.
- More importantly, even if one had collision detection and sensed no collision when sending, a collision could still occur at the receiver. This situation results from the particular characteristics of the wireless channel.

Given the difficulties with detecting collisions at a wireless receiver, the designers of IEEE 802.11 developed an access protocol that aimed to avoid collisions (hence, the name CSMA/CA), rather than detect and recover from collisions (CSMA/CD).

First, the IEEE 802.11 frame contains a duration field in which the sending station explicitly indicates the length of time that its frame will be transmitting on the channel. This value enables other stations to determine the minimum amount of time (the so-called network allocation vector, [NAV]) for which they should defer their access.

The MAC layer specification for 802.11 has similarities to the 802.3 Ethernet wired line standard. The protocols for 802.11 use a protocol scheme know as carrier sense multiple access collision avoidance (CSMA/CA). This protocol avoids collisions instead of detecting a collision like the algorithm used in 802.3. It is difficult to detect collisions in an RF transmission network, which is why collision avoidance is used. The MAC layer operates together with the physical layer by sampling the energy over the medium transmitting data. The physical layer uses a clear channel assessment (CCA) algorithm to determine if the channel is clear. This is accomplished by measuring the RF energy at the antenna and determining the strength

of the received signal. The best method to use depends upon the levels of interference in the operating environment. The CSMA/CA protocol provides options that can minimize collisions by using RTS, CTS, data and acknowledge (ACK) transmission frames, in a sequential fashion.

Communications is established when one of the wireless nodes sends a short message RTS frame. The RTS frame includes the destination and the length of message. The message duration is known as the *network allocation vector* (NAV). The NAV alerts all others in the medium to back off for the duration of the transmission. The receiving station issues a CTS frame, which echoes the sender's address and the NAV. If the CTS frame is not received, it is assumed that a collision occurred and the RTS process starts over. After the data frame is received, an ACK frame is sent back verifying a successful data transmission. A common limitation with wireless LAN systems is the hidden node problem. This can disrupt 40 percent or more of the communications in a highly loaded LAN environment. It occurs when a station in a service set cannot detect the transmission of another station to detect that the media is busy. Figure 21-18 shows the hidden node.

With the emphasis now on Voice over IP telephone applications, the 802.11 standard enables the interconnection of the traditional PBX and LANs as shown in Figure 21-19.

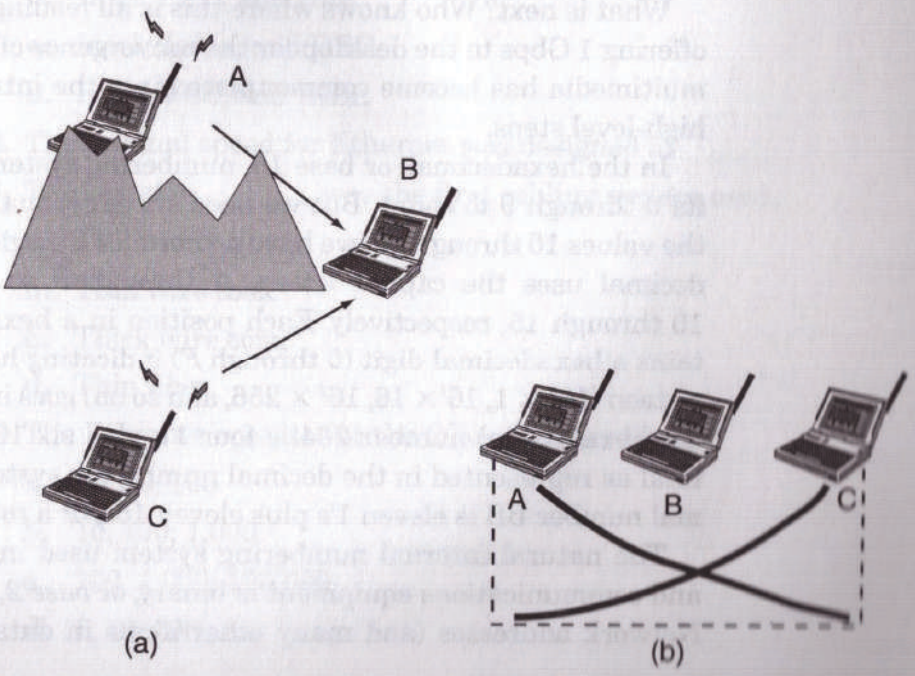


Figure 21-18
The hidden node

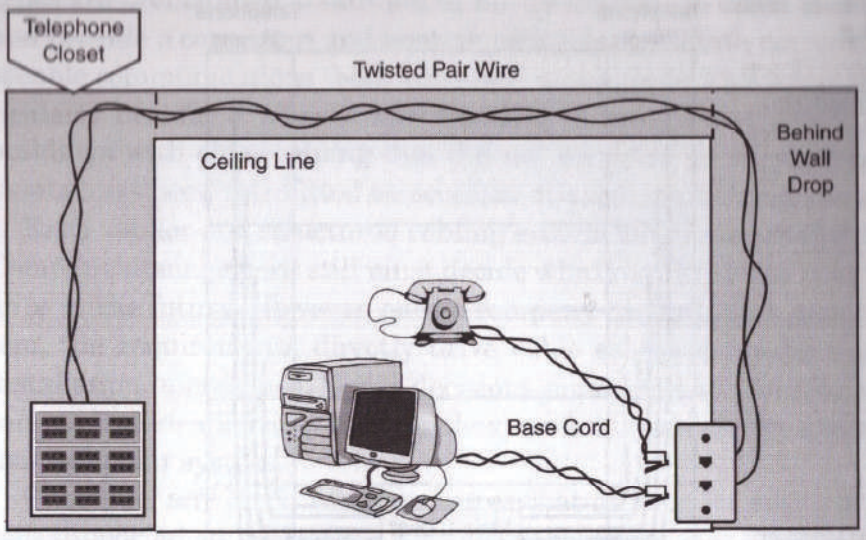


Figure 39-3
The components of the horizontal distribution

Twisted Pair

Twisted pair refers to telephone wiring. This type of paired wiring usually comes as two solid copper wires, each covered with a plastic insulator and twisted together. It is not the silver satin cable often sold in hardware stores for use with telephones. Silver satin is a flat-wire system where all the wires are laid side by side rather than twisted around each other in pairs. Twisted-pair wiring is often described as shielded or unshielded twisted pair (STP or UTP). UTP (Figure 39-4) has been primarily used for voice communications. It is twisted to reduce induction effects in adjacent wires, otherwise known as cross talk. The more twists per foot, the better the reduction in the cross talk. The twisted-pair wiring is usually sold in various flavors, as shown in Table 39-1. In these pairings, the number of pairs is always important because of basic and immediate needs as well as future needs. The best approach anyone can take when installing wiring to the desk or in vertical distribution is always to add more.



Figure 39-4
The unshielded twisted-pair wiring

Table 39-1
Typical Number of Pairs Used for Installation

Description	Number of Pairs	Used
Analog telephone wiring	2 pairs	1 pair
Digital telephone wiring	1–4 pairs	1–4 pairs
LAN wiring Ethernet 10BASE-T	4 pairs	2 pairs
LAN wiring token ring 4 or 16 Mbps	4 pairs	2 pairs
Horizontal telephone wires	25 pairs	1–4 for individual sets

The wire pairs are typically color coded so that the pairing can be kept together visually, rather than by guesswork. This pairing color code is designed so that a striped wire and a solid-color wire, or a striped wire and a striped wire of the reverse color combination, create a pair; older wiring systems, such as the old telephone wiring in households, use a solid-color wire convention. Table 39-2 shows the more common combinations. The two colors on the wire insulation represent where the convention uses a striped color code (that is, white/blue means a white wire with a blue stripe,

Table 39-2
Color Coding for Normal Pairings

Pair Number	Color-Coded Pairing
One	White wire/blue stripe and blue wire, or white wire/blue stripe and blue wire/white stripe
Two	White wire/orange stripe and orange wire, or white/orange and orange/white
Three	White/green and green or white/green and green/white
Four	White/brown and brown or white/brown and brown/white
Older analog wiring found still in some buildings uses a four-wire pairing with these color-coding conventions.	
One	Red/green Red is the transmit and green is the receive wire.
Two	Black/yellow Black is the battery and yellow is the ground for some sets; if a two-line set is used, then black is transmit 2 and yellow is receive 2.

blue/white means a blue wire with a white stripe). In the reverse wiring, the color code uses the two colors in the reverse; however, the normal pair contains a striped and a solid wire (that is, white/blue is a white wire with blue stripes and blue is a solid blue wire).

If installing a telephone set (for example) that requires one pair, it would be imprudent to run the wires to the set with only one pair enclosed. Otherwise, the first thing that would happen if you chose to change the device hanging off these wires is that you would be confronted with the need to rewire with additional pairs. Telephone systems typically can be used that require from one to four pairs at the desk. Thus, the minimum number of pairs to the desk is four. We have always adopted a standard that a minimum of two each of the four-pair wiring be installed. After all, the cost of adding a few hundred feet of wire at the time of initial installation is only a few cents per foot. If, however, rewiring is necessary, the labor to go back and pull more wires will cost a significant amount (Table 39-3).

UTP is a balanced wiring system. This means that both wires carry similar types of electrical signals and have similar characteristics. No special grounding is required. With the appropriate adapter, UTP can be used to support nonremote terminal traffic that would otherwise require RS232 cabling. Shielded twisted-pair (STP) wiring (Figure 39-5) was endorsed by IBM at about the same time the company released the 4-Mbps token ring.

Table 39-3

Comparing the Difference Between Pulling a Single Cable Twice versus Two Cables Once Shows the Real Dollar Value of Preplanning. The Main Cost Difference Is the Labor.

Item	Cost of Pulling One Cable to Each Desk (each)	Cost of Pulling Two Cables to Each Desk (each × 2)
Materials (4-pair cable, block, RJ45, cross-connect wire, base cord)	$9.35	$18.70
Labor to pull wires, punch down and test	$300	$320
Total cost per run	$309.35	$338.70
@ 50 pulls	$15,467.50	$16,935
Cost for additional installations	$15,467.50	$0 50
Total cost	$30,935	$16,935
Difference		($14,000)

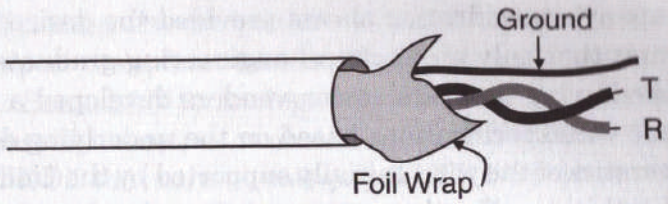


Figure 39-5
Shielded twisted pairs include foil wrap that is connected to a good earth ground.

STP was better for data communications than early UTP was, but only if it was installed correctly. The shield—a foil sheath covering the copper pairs—must be properly grounded. If not, the shield can create a ground loop (Figure 39-6). This situation introduces the possibility that a wiring system will become an antenna, which is worse than if no shielding were used. STP is also more costly than UTP because of the extra foil shielding and because the jacks and plugs cost significantly more than those used with UTP. STP was generally used only for token ring installations, before appropriate UTP grades and components became available. The shield, if properly installed, will provide better cable performance, such as reduced cross talk, protection from grounds and faults, greater signal distances, and protection from EMI and RFI interference.[2]

AT&T (now Lucent) and others experienced many queries from customers as to the transmission characteristics of various UTP offerings. The

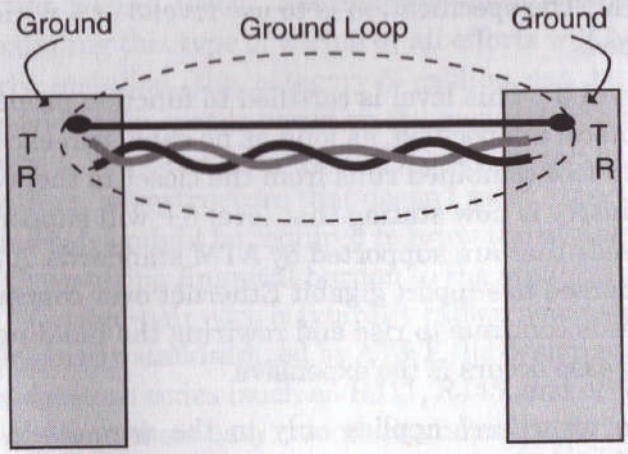


Figure 39-6
If grounded on both ends at different potentials, the shield creates a ground loop. This is worse than no ground at all.

[2] EMI refers to electromechanical interference; RFI refers to radio frequency interference. These are the conditions that are introduced into the copper on the twisted-pair wiring.

company's specification sheets provided the desired information, but in a format that only an electrical engineering graduate could understand. To make buying products easier, vendors developed a set of simplified data-grade wire specifications based on the underlying detailed electrical characteristics of the wire. Initially supported by the Underwriters Laboratories (UL), this specification contained three levels numbered 1 through 3. As more wiring products became available and needed, the company extended the grading scheme. Now five defined levels or data grades for UTP wiring are supported by the TIA/EIA, each capable of handling its own requirements as well as those addressed by all lower levels:

- **Level 1** Sometimes described as barbed wire, this level will handle the requirements of voice communications, called plain old telephone service (POTS), and limited data rates that meet EIA 232-specific data rates (such as 19.2 Kbps).
- **Level 2** This level is certified to handle the requirements of a 4-Mbps token ring.
- **Level 3** This level addresses Ethernet's requirements of 10 Mbps on twisted pairs. The specifications from various vendors suggest that this specification will support higher rates, such as 16 and 20 Mbps, but it is not an EIA standard.
- **Level 4** This level is used for 16-Mbps token ring speeds and the 20-Mbps ARCNet service. IBM is planning to introduce a 25-Mbps token-passing ring concept at duplex arrangements that will be called ATM. The specification is to use level 3 or 4 wiring to support these speeds.
- **Level 5** This level is certified to function properly at up to 100 Mbps, protocol unspecified, as long as no cable run exceeds 100 meters (328-foot combined runs from the closet to the workstation). The industry is now stating that level 5+ will support up to 155-Mbps speeds that are supported by ATM standards. A level 5E is being endorsed to support gigabit Ethernet over copper cables because the speeds continue to rise and rewiring the building every time a speed increase occurs is too expensive.

The term *level* applies only to the wire itself. The entire system is described as category *n*, where n matches the wiring level. The entire system includes jacks, plugs, and patch panels, in addition to the wiring itself. These components also need to be certified as meeting category whatever requirements. Because category 5 components cost more to install, people sometimes install category 5 wiring with category 3 components. Because

replacing the components is much easier than rewiring a building, they expect to replace the components later if required.

Some LAN technologies require 100 Mbps. These include two varieties of Fast Ethernet, as well as a standard for Fiber Distributed Data Interface (FDDI) that runs on category 5 UTP. No matter what LAN topology is used, UTP is almost always configured with a radial wiring plan and rarely as a physical bus. Radial (or star, radiating out from a central point) fits well into building cabling troughs, which in large part accounts for the popularity of UTP data wiring.

UTP offers the lowest medium cost per foot installed. Of available technologies, it is the most limited in distance because of the nature of the medium. Every medium has an inverse relationship between distance and potential speed. The specific ratio varies with the media. For any given speed, UTP has the shortest distance limitation of the media discussed in this chapter.

However, UTP has many suppliers, and installers are familiar with it. Because of its popularity, vendors focus on getting their products to function in this type of physical plant. Unless a company will possibly need speeds greater than 100 to 155 Mbps to the desktop, or will need cable runs to exceed 100 meters, a category-5 installation is likely to handle all requirements for at least the next 5 years. If you are planning a category-5 implementation, however, be sure to have installers who are experienced with category 5 and who have more than just basic telephone wiring experience. Installing category 5 is not particularly difficult, but it is different from the other wiring levels. Specific considerations must be adhered to when installing this type of wiring or all efforts will be for naught. When improperly installed, the category-5 cabling can be negated and functionally provide only speeds based on the weakest link in the system, such as category 1. You can spend a lot of money on the wiring system, only to be left with an infrastructure that doesn't meet expectations; or, retrofits that are expensive might be required to bring the wiring up to standard. This can be a devastating financial burden to the organization.

Twisted-pair wire is typically cabled using standardized jacks and plugs. Originally standardized by AT&T, the designations of these jacks are easily understood codes (such as RJ11, RJ45, and others). We assume that someone found them easy to understand, anyway. In any case, the key differentiating factors among these jacks and plugs are

- How many pairs of wire can terminate on the jack (two, three, or four)?
- How many pairs of wire actually do terminate on the jack?
- How are the pairs assigned?

Every electrical signal path on UTP requires one pair of wires. If two signal paths are required (for example, digital sending and receiving), then two pairs are required. However, those pairs most likely will not be assigned in sequence (for example, left to right or right to left); that would be too easy. As an example, voice wiring on an RJ11 jack (which can hold two pairs of wire and thus handle two analog voice lines) typically uses the center pair of pins or wires on the jack for the first line, and the outer pair for the second line if present. If only one line is to be installed, sometimes the outer pair will not even be present on the jacks and plugs, even though the RJ11 jacks are still used.

Every defined use of UTP (for example, 10BASE-T for Ethernet) includes a very precise specification of which positions will be used on what jack for which transmission path. 10BASE-T is different from token ring, and both differ from the standard for FDDI[3] over UTP, even though all of these are defined on RJ45 jacks and plugs. See Figure 39-7 for a representation of the various ways that the wires can be used in the jacks based on differing standards.

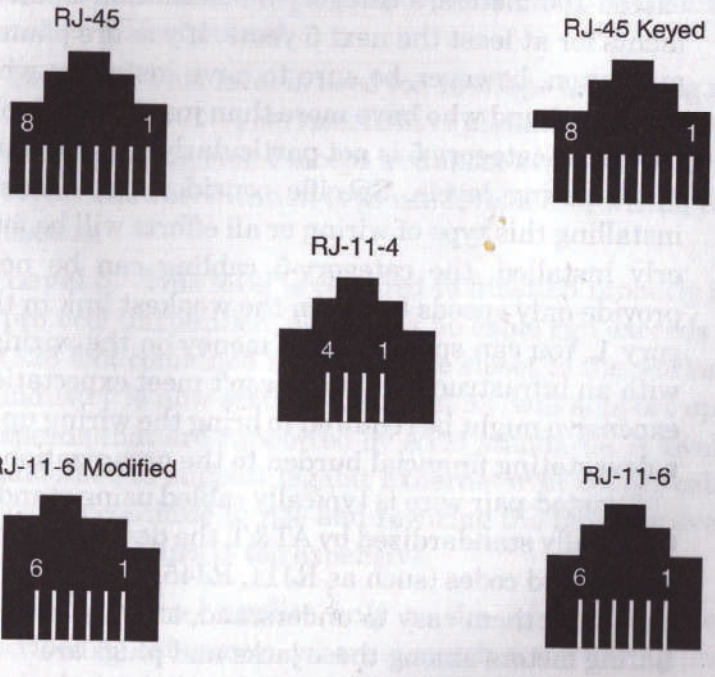


Figure 39-7
The specific number of pairs and the pinouts differ on the basis of use.

[3] Fiber Distributed Data Interface (FDDI) was specified on a proprietary cabling scheme using Copper Distributed Data Interface (CDDI). This has not been popular because of the Ethernet speed increases and proliferation.

Coax

Coaxial cable is decreasing in popularity. Before fiber was developed, it was the only practical medium to achieve high bandwidths over distances exceeding a few hundred feet. Whereas a high-speed signal (up to about 100 Mbps) can go via UTP up to 100 meters or so, some forms of coaxial cable can carry signals requiring several hundred Mbps a few thousand meters. The two basic forms of coaxial cable are baseband and broadband. Baseband is used for an Ethernet connection and broadband is the typical cable used in cable TV (CATV) systems. The use of baseband versus broadband systems was covered in Chapter 23.

Coaxial cable is an unbalanced medium; the signal goes through one conductor (the inner core), while the other conductor (the outer braid) functions as the ground. Signals can be converted from a balanced to an unbalanced medium (and back) via use of a special passive device called a balun (balanced and unbalanced), so long as the numbers of required conductors match. This technique is used to supplement coaxial cable with twisted-pair wires as an economic measure. Coaxial ends might be used on a cable run, but UTP wires are used in the middle for the vertical or horizontal run to save money in the wiring of a building. See Figure 39-8 for a typical installation using the balun.

All coaxial cables have the same basic construction. Each has an inner conductive core, usually of copper (usually solid but sometimes stranded). Around the core is the dielectric, an insulating material. Around the dielectric is a second conductive path or sheath, usually in the form of braided copper. This basic construction is illustrated in Figure 39-9 for baseband coaxial cable and in Figure 39-10 for broadband coaxial cable. As with other available media, the underlying construction is varied in many ways to enable customization for specific purposes.

Variations in the thickness of the three key components, as well as in the type or types of covering used, result in a bewildering array of possible cables. Each specific type has a designation (for example, RG58A) that

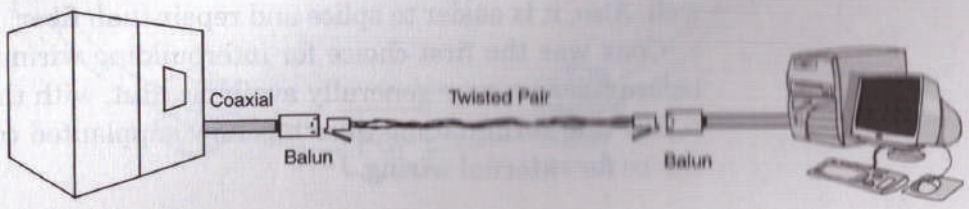


Figure 39-8
The balun is used to run less-expensive twisted pairs to a coax interface.

Do Not Remove This Wrapper

University of Pittsburgh
University Library System

DATE DUE

JAN - 2 2002			
JUL - 1 2002		MAY - 9 2003	
AUG 15 2002			
JUN 24 2003			
GAYLORD			PRINTED IN U.S.A.

APPENDIX JACOBSEN04

MARC display

Save this record | Return to record

```
LEADER 00000nam a22000003a 4500
005    19950119000000.0
008    950119s1994    dcu           10    eng r
020    0940272229
040    CaOON|beng|cCaOON
090    TK5101.A1 N277 1994
110 2  National Cable Television Association.|bConvention and
       Exposition|n(43rd :|d1994 :|cNew Orleans, La.)
245 10 National Cable Television Association, 1994 technical
       papers :|b43rd Annual NCTA Convention and Exposition, New
       Orleans, Louisiana, May 22-25, 1994.
260    [Washington, D.C.] :|bThe Association,|c[c1994]
300    484 p. :|bill. ;|c28 cm.
500    Spine title: 1994 NCTA technical papers, the National Show,
       New Orleans, Louisiana, May 22-25, 1994.
500    Compiled by Katherine Rutkowski.
650  0 Telecommunication|xCongresses.
650  0 Telecommunication policy|xCongresses.
650  0 Cable television|zUnited States|xCongresses.
650  0 High-definition television|xCongresses.
700 1  Rutkowski, Katherine S.
710 2  National Cable Television Association.
```

Location	Call number & Local holdings	Status
NRC Ottawa	TK5101.A1 N277 1994	Available

Date modified: 2024-07-11

Français

Government of Canada Gouvernement du Canada

Search Canada.ca

Canada.ca › Science and innovation › Federal Science Libraries Network › Catalogue › Search results

Record number b1596006 Search

Save this record MARC display

Corporate author	National Cable Television Association. Convention and Exposition (43rd : 1994 : New Orleans, La.)
Title	**National Cable Television Association, 1994 technical papers : 43rd Annual NCTA Convention and Exposition, New Orleans, Louisiana, May 22-25, 1994.**
Publisher	[Washington, D.C.] : The Association, [c1994]

Request this item

Location	Call number & Local holdings	Status
NRC Ottawa	TK5101.A1 N277 1994	Available

ISBN	0940272229
Description	484 p. : ill. ; 28 cm.
Note	Spine title: 1994 NCTA technical papers, the National Show, New Orleans, Louisiana, May 22-25, 1994.
	Compiled by Katherine Rutkowski.
Other author	Rutkowski, Katherine S.
Corporate name	National Cable Television Association.
Subject	Telecommunication -- Congresses.
	Telecommunication policy -- Congresses.
	Cable television -- United States -- Congresses.
	High-definition television -- Congresses.

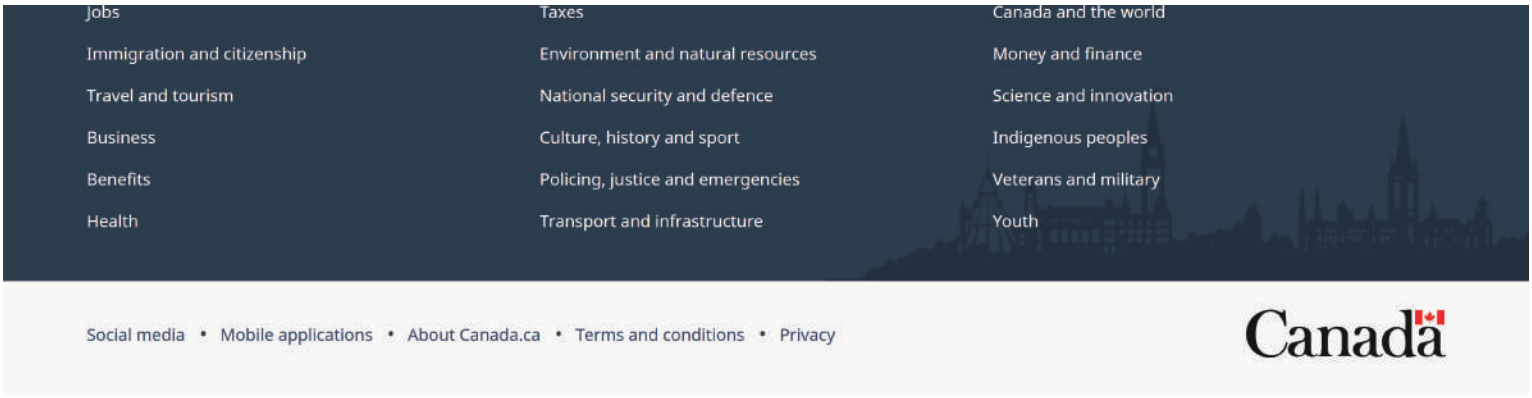
Permanent link

Save this record MARC display

Date modified: 2024-07-11

Government of Canada

All Contacts Departments and agencies About government


Jobs
Immigration and citizenship
Travel and tourism
Business
Benefits
Health
Taxes
Environment and natural resources
National security and defence
Culture, history and sport
Policing, justice and emergencies
Transport and infrastructure
Canada and the world
Money and finance
Science and innovation
Indigenous peoples
Veterans and military
Youth

APPENDIX JACOBSEN05

IEEE.org IEEE *Xplore* IEEE SA IEEE Spectrum More Sites Donate Cart Create Account Personal Sign In

Browse My Settings Help Access provided by: Penn State University Sign Out

Access provided by: Penn State University Sign Out

All

ADVANCED SEARCH

Conferences > Proceedings of GLOBECOM '95

Synchronized DMT for multipoint-to-point communications on HFC networks

Publisher: IEEE Cite This PDF

<< Results

K.S. Jacobsen ; J.A.C. Bingham ; J.M. Cioffi All Authors

8 Cites in Papers

16 Cites in Patents

89 Full Text Views

Alerts

Manage Content Alerts

Add to Citation Alerts

Abstract

Authors

References

Citations

Keywords

Metrics

More Like This

Downl PDF

Abstract:
This paper presents a method for synchronizing a large number of discrete-multitone-based modems for multipoint-to-point communications on hybrid fiber-coax (HFC) network... **View more**

Metadata

Abstract:
This paper presents a method for synchronizing a large number of discrete-multitone-based modems for multipoint-to-point communications on hybrid fiber-coax (HFC) networks. Each newly-installed remote unit is first loop-timed to the central modem master clock. Ranging is then performed to determine the sample delay required to align at the central unit receiver the symbols transmitted by the remote with those transmitted by remote units already operating. Upon completing the synchronization procedure, the remote unit transmits a signal to train the central unit receiver so that the available upstream bandwidth can be used as efficiently as possible during steady-state operation.

Published in: Proceedings of GLOBECOM '95

Date of Conference: 14-16 November 1995

DOI: 10.1109/GLOCOM.1995.502547

Date Added to IEEE *Xplore*: 06 August 2002

Publisher: IEEE

Print ISBN:0-7803-2509-5 **Conference Location:** Singapore

References is not available for this document.

Authors

References

Download PDFs | Export | References & Cited By

Select All

1. K. S. Jacobsen and J. M. Cioffi, "An Efficient Digital Modulation Scheme for Multimedia Transmission on the Cable Television Network", *43rd Annual National Cable Television Association (NCTA) Convention and Exposition*, 1994-May.
   Google Scholar

2. J. A. C. Bingham and K. S. Jacobsen, "CATV Reverse Channel Transmission Using Synchronized DMT Part 1. Overview: Network Data Rates and Services Protocol", *IEEE 802.14 Working Group Contribution number 95-001*, 1995-January.
   Google Scholar

3. K. S. Jacobsen, J. A. C. Bingham and J. M. Cioffi, "CATV Reverse Channel Transmission Using Synchronized DMT Part 2: Synchronization and Training", *IEEE 802.14 Working Group Contribution number 95-002*, 1995-January.
   Google Scholar

4. J. A.C. Bingham and K. S. Jacobsen, "CATV Reverse Channel Transmission Using Synchronized DMT Part 3: Data Access Protocol", *IEEE 802.14 Working Group Contribution number 95-003*, 1995-January.
   Google Scholar

5. J. A.C. Bingham, "Multicarrier Modulation for Data Transmission: An Idea whose Time Has Come", *IEEE Communications Magazine*, vol. 28, no. 5, pp. 5-14, May 1990.
   View Article Google Scholar

6. J. M. Cioffi, "A Multicarrier Primer", *ANSI T1E1.4 Committee Contribution 91-157*, 1991-November.
   Google Scholar

Citations

Keywords

Metrics

Back to Results

More Like This

Synchronized discrete multi-tone (SDMT) modulation for cable modems: making the most of the scarce reverse channel bandwidth
WESCON/97 Conference Proceedings
Published: 1997

The evolution of Hybrid Fiber-Coaxial cable networks to an all-fiber network
2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference
Published: 2011

Show More

IEEE Personal Account
CHANGE USERNAME/PASSWORD

Purchase Details
PAYMENT OPTIONS
VIEW PURCHASED DOCUMENTS

Profile Information
COMMUNICATIONS PREFERENCES
PROFESSION AND EDUCATION
TECHNICAL INTERESTS

Need Help?
US & CANADA: +1 800 678 4333
WORLDWIDE: +1 732 981 0060
CONTACT & SUPPORT

Follow

About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | IEEE Ethics Reporting | Sitemap | IEEE Privacy Policy

A public charity, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.

© Copyright 2025 IEEE - All rights reserved, including rights for text and data mining and training of artificial intelligence and similar technologies.

IEEE Account

» Change Username/Password

» Update Address

Purchase Details

» Payment Options

» Order History

» View Purchased Documents

Profile Information

» Communications Preferences

» Profession and Education

» Technical Interests

Need Help?

» **US & Canada:** +1 800 678 4333

» **Worldwide:** +1 732 981 0060

» Contact & Support

About IEEE *Xplore* | Contact Us | Help | Accessibility | Terms of Use | Nondiscrimination Policy | Sitemap | Privacy & Opting Out of Cookies

A not-for-profit organization, IEEE is the world's largest technical professional organization dedicated to advancing technology for the benefit of humanity.
© Copyright 2025 IEEE - All rights reserved. Use of this web site signifies your agreement to the terms and conditions.

Synchronized DMT for Multipoint-to-point Communications on HFC Networks

Krista S. Jacobsen, John A.C. Bingham*, and John M. Cioffi
Information Systems Laboratory, Stanford University, Stanford, CA 94305-4055
Phone: (415) 723-2525 FAX: (415) 723-8473
e-mail: jacobsen@isl.stanford.edu

Abstract

This paper presents a method for synchronizing a large number of discrete-multitone-based modems for multipoint-to-point communications on hybrid fiber-coax (HFC) networks. Each newly-installed remote unit is first loop-timed to the central modem master clock. Ranging is then performed to determine the sample delay required to align at the central unit receiver the symbols transmitted by the remote with those transmitted by remote units already operating. Upon completing the synchronization procedure, the remote unit transmits a signal to train the central unit receiver so that the available upstream bandwidth can be used as efficiently as possible during steady-state operation.

1 Introduction

In the near future, the deployment of hybrid fiber-coax (HFC) networks by both cable television and telephone companies will enable high-speed digital transmission both to and from customer premises. Generally, the network topology will be a tree-and-branch configuration, with fiber emanating from a central site, as shown in Figure 1. In this configuration, the fiber and a significant percentage of the coaxial cable may serve hundreds of users. For either point-to-point or point-to-multipoint communications in the downstream direction (from the central site to the remote units), synchronization of the remote units is not of importance because the central unit coordinates all downstream communications. In the upstream direction, however, the remote units must be synchronized to ensure their transmissions do not interfere with each other.

Discrete multitone (DMT) is a powerful modulation technique that can optimize the data rate transmitted on a channel. It has been shown in [1] that, with reasonable computational complexity, DMT can achieve high data rates on channels with a variety of transmission impairments, including passband ripple, spectral nulls, and radio-frequency interference, all of which can be expected in HFC networks. This paper presents a procedure for synchronizing a large number of DMT-based remote units for upstream communications so that the advantages of DMT, particularly its ability to avoid narrowband interference while still supporting high data rates, can be exploited in multipoint-to-point HFC environments.

Section 2 explains the operation of the DMT transceiver. The synchronization requirements imposed by multipoint-to-point environments are presented in Section 3, along with procedures for synchronizing and training multiple DMT-based remote units. Finally, Section 4 concludes that the synchronization procedure presented enables efficient multipoint-to-point communications on HFC networks.

2 Discrete Multitone Modulation

Discrete multitone modulation uses an inverse discrete Fourier transform (IDFT) to partition a transmission channel into a set of orthogonal, equal-bandwidth subchannels such that the frequency response of each subchannel is approximately constant across its bandwidth, as illustrated in Figure 2. The resulting subchannels are essentially memoryless, and the small amount of interblock and intersubchannel interference caused by the channel can be eliminated by the use of a cyclic prefix. Thus, a traditional equalizer is not required. In the transmitter, a bitstream is encoded as a set of quadrature amplitude modulated (QAM) subsymbols. Each QAM subsymbol represents a number of bits determined by the signal-to-noise ratio (SNR) of its associated subchannel. As a simple example, assume a transmission channel is divided into 4 subchannels[1] with SNRs of 13.5, 22.5, 17.0, and 6.0 dB, respectively. Then, for a received (uncoded) symbol error probability of 10^{-7}, the subchannels can support 2, 5, 3, and 0 bits, respectively. Thus, DMT can transmit 10 bits per symbol with this configuration, with the first three subchannels supporting 4-QAM (or, equivalently, QPSK), 32-cross QAM, and 8-cross QAM. The total bit rate of the system is then (10 bits/symbol)(W symbols/s), where W

*John Bingham is with Amati Communications Corporation in Mountain View, CA.

[1] In reality, four subchannels would rarely be sufficient to satisfy the requirement of flat frequency response magnitude over the bandwidth of each subchannel.

0-7803-2509-5/95 US$4.00 © 1995 IEEE

Authorized licensed use limited to: Penn State University. Downloaded on January 07,2025 at 15:58:48 UTC from IEEE Xplore. Restrictions apply.

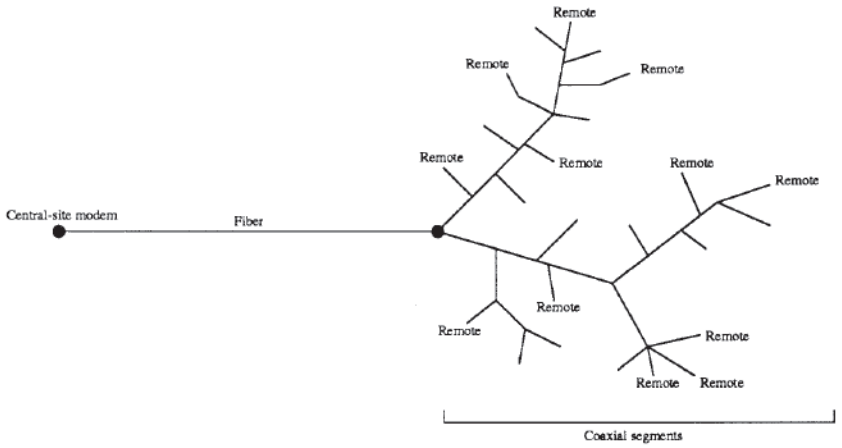


Figure 1: HFC network configuration

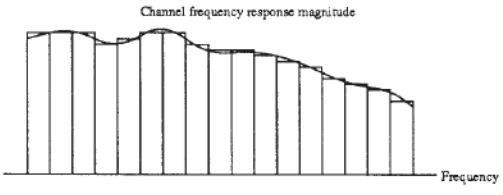


Figure 2: A channel partitioned into subchannels

is the nominal bandwidth of each subchannel. Because the SNR of the last subchannel is not high enough to support 2 bits at the desired error probability, it is not used. The constellation points are input to a complex-to-real IDFT, and the output is converted from digital to analog format. The signal, a sum of real time-domain sinusoids, is then upconverted to the desired carrier frequency and transmitted over the channel. In the receiver, after downconversion and analog-to-digital conversion, the data are input to a DFT. Each output value is then scaled by a single complex number to compensate for the magnitude and phase of each subchannel's response, and a memoryless detector decodes the resulting symbols. The set of complex multipliers is known as the frequency-domain equalizer (FEQ). Changes in the channel magnitude and/or phase are accommodated by updating the FEQ taps as the system operates. A generalized DMT transmitter and receiver pair is shown in Figure 3.

During steady-state operation, the subchannel SNRs are monitored, and the bit distribution is modified as necessary to maintain optimal system performance. This adaptivity enables DMT to alleviate problems caused by changes in the channel frequency response or noise while the system is in use. If the attenuation or noise of one or more subchannels becomes severe enough to compromise the system error performance (which could happen, for example, if a strong external radio-frequency signal suddenly coupled into the network), the transmitter simply assigns fewer or no bits to that subchannel while the noise or attenuation persists. Therefore, even unpredictable, bursty, high-power noise such as ingress from amateur radio or other radio-frequency interferers can be tolerated by a multicarrier system, assuming that either the resulting channel can continue to support the required data rate or the data rate can be decreased. Dramatic changes in the channel itself can also be accommodated similarly by an adaptive DMT system, provided that the adaptation speeds of the DMT algorithms are sufficient to track the changes. Furthermore, changes in the channel frequency response or noise during operation that actually improve any of the subchannel SNRs can be used to increase the system's overall data rate, if desired.

3 Synchronized DMT for Multipoint-to-point Communications

At all times, including during installation, communications from any particular remote unit to the central-site modem in a multipoint-to-point HFC environment must occur without disrupting in-progress data transmissions by other remote units. Thus, when it is first installed on the network, before a remote unit is synchronized it may only transmit while the other remote units are silent. To accommodate this requirement, silent intervals of a predetermined length are observed periodically by all remote units on the network. The duration of these silent periods must be such that an installation signal transmitted by any remote on the network will arrive at the central-site modem before the end of the silent interval. Specifically, the silent interval must be at least the maximum variation

Authorized licensed use limited to: Penn State University. Downloaded on January 07,2025 at 15:58:48 UTC from IEEE Xplore. Restrictions apply.

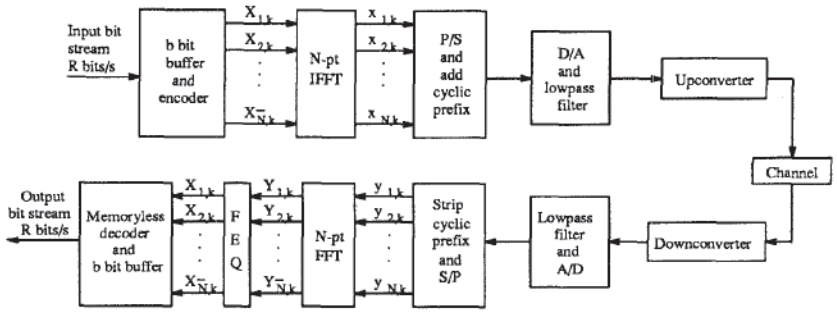


Figure 3: DMT block diagram

in round-trip signal delay over all remotes plus the time required to transmit the information, which may be only one symbol period.

In a point-to-point environment, use of the IDFT with a cyclic prefix added to each symbol ensures that the subchannels are orthogonal in the presence of white Gaussian noise. In a multipoint-to-point environment, however, there are two additional requirements for achieving independent subchannels. First, only one remote unit may use any given subchannel during each DMT symbol period.[2] Second, the remote units must be synchronized such that their transmissions arrive at the central-site unit at the appropriate times. Because the distances between remote units and the central unit will in general differ significantly, the times required for signals to travel from each remote unit to the central-site modem will also vary. Therefore, in addition to synchronizing with the central-site modem clock, each remote unit must acquire its sample delay before it can begin operating. The minimum required delay for a given remote is the difference between the maximum round-trip signal delay from the central-site modem to all units on the network and the remote unit's round-trip signal delay. In general, the maximum round-trip signal delay will correspond to the unit furthest from the central site.

3.1 Synchronization Procedure

When a remote unit is installed on a network, it must first loop-time its local clock with the central-site modem clock. The central unit transmitter then sends a signal in a downstream control channel that instructs the remote unit to send its installation parameters, which include its factory-assigned serial number so that the central-site modem can address transmissions during installation to the installing remote in the downstream control channel. Because the bit capacities of the subchannels from this remote to the central site have not yet been determined, the remote unit uses differential QPSK (2 bits per subchannel) to transmit its serial number and other required parameters on some subset (or perhaps all) of the available subchannels during the next available quiet period. When the central unit controller detects this signal, it compares the signal's symbol boundaries to those of symbols transmitted by the remote units currently operating. In general, there will be a difference in the symbol boundaries, and the controller computes and sends the time delay required of the synchronizing unit to correct the misalignment. The remote unit then implements the requested sample delay and transmits a signal requesting verification that it is synchronized. If the remote unit transmission is indeed synchronous, the central unit controller sends a signal to that unit in the downstream channel to indicate that no further shifting is required, and that the remote unit may now communicate with the central-site modem incorporating the appropriate delay. Otherwise, the synchronization procedure is repeated until the central-site controller determines the remote unit is synchronized. After the initial symbol delay has been determined, unless a remote unit is moved or its connection to the network is terminated, it should not have to be resynchronized. Failing to synchronize the remote units to within a certain tolerance can result in interchannel interference, which can decrease the achievable bit rates on the affected subchannels.

3.2 Training

After receiving its installation parameters from the central-site modem and incorporating the required sample delay, the remote unit transmits a "wide-band" signal during a specified number of upcoming silent periods to train the central unit receiver.[3] Because the newly installed remote unit is now synchronized with respect to the other remote

[2] Note that the subchannels assigned to a particular remote unit need not be the same from symbol to symbol. For low bit rate applications, for example, sufficient throughput may be achieved by using one subchannel (but not necessarily the *same* subchannel) every few symbols.

[3] The term "wide-band" here means that the remote unit is allowed to use all of the allowed subchannels in its bandwidth during a specified number of upcoming silent periods. Note that some of the subchannels may be disallowed for some reason. For example, if the network operator has another system in place on the network using bandwidth that overlaps the bandwidth of the DMT system, then the overlapping DMT subchannels are marked as "taboo", and they are never used for either training or data transmission.

Authorized licensed use limited to: Penn State University. Downloaded on January 07,2025 at 15:58:48 UTC from IEEE Xplore. Restrictions apply.

units, it can transmit using all of the symbols during the next several silent periods for channel analysis. All other remote units remain quiet while the remote unit transmits a training signal on the permissible subset of the subchannels allocated to it, and the central unit controller records the bit capacity and FEQ tap setting of each subchannel from that remote unit. The bit capacities are used to determine subchannel assignments when the remote later requests either a constant data rate or a packet transmission. Because the controller allocates the subchannels to the various remote units every symbol period, it can apply the appropriate FEQ tap to each subchannel during demodulation. Thus, if the remotes are all properly synchronized, the signal arriving at the central unit receiver, which is actually an aggregate of transmissions from a number of different remote units, can be demodulated as though it were from a single remote modem, using the appropriate mixture of FEQ tap settings.

After a remote has been installed, it is periodically retrained during another silent interval reserved specifically for this purpose. As during the installation silent period, all remote units that are not training remain quiet to allow the central unit controller to update its bit capacities and FEQ tap settings for the training remote. Depending on the frequency of these silent intervals, the number of remotes on a particular network, and other system parameters, each remote could be retrained as often as many times per second or as infrequently as every two (or more) seconds.

4 Conclusion

In this paper, a procedure for installing DMT-based remote units in a multipoint-to-point HFC environment has been described. Specifically, steps and requirements for synchronizing and training remote units have been outlined. Although traditionally DMT has been used only in point-to-point environments, this paper shows that DMT can also be employed effectively in multipoint-to-point environments, as long as the remote units are synchronized to within a certain tolerance. Thus, the advantages of DMT, specifically its efficient use of the available bandwidth and its robustness to radio-frequency ingress and other noises, can be exploited for transmission on multipoint-to-point HFC networks.

References

[1] K.S. Jacobsen and J.M. Cioffi. "An Efficient Digital Modulation Scheme for Multimedia Transmission on the Cable Television Network". In *Technical Papers, 43rd Annual National Cable Television Association (NCTA) Convention and Exposition*, New Orleans, LA, May 1994.

[2] J.A.C. Bingham and K.S. Jacobsen. "CATV Reverse Channel Transmission Using Synchronized DMT, Part 1. Overview: Network, Data Rates and Services, Protocol". In *IEEE 802.14 Working Group Contribution number 95-001*, Boston, MA, January 1995.

[3] K.S. Jacobsen, J.A.C. Bingham, and J.M. Cioffi. "CATV Reverse Channel Transmission Using Synchronized DMT, Part 2: Synchronization and Training". In *IEEE 802.14 Working Group Contribution number 95-002*, Boston, MA, January 1995.

[4] J.A.C. Bingham and K.S. Jacobsen. "CATV Reverse Channel Transmission Using Synchronized DMT, Part 3: Data Access Protocol". In *IEEE 802.14 Working Group Contribution number 95-003*, Boston, MA, January 1995.

[5] J.A.C. Bingham. "Multicarrier Modulation for Data Transmission : An Idea whose Time Has Come". *IEEE Communications Magazine*, 28(5):5–14, May 1990.

[6] J.M. Cioffi. "A Multicarrier Primer". In *ANSI T1E1.4 Committee Contribution 91-157*, Boca Raton, FL, November 1991.

Authorized licensed use limited to: Penn State University. Downloaded on January 07,2025 at 15:58:48 UTC from IEEE Xplore. Restrictions apply.

APPENDIX LAWTON01

PennState University Libraries | Catalog

Bookmarks 0 Course Reserves My Account

Keyword Search... Search Advanced search Start Over

Share Bookmark Report an Issue

MARC View

LEADER 11231nas a22012971 4500			
001 156262			
003 SIRSI			
005 20151212044536.0			
008 820329c19669999xx mu 0 0eng d			
010			a| 74648480
019			a| MARS
022			a| 0018-9162
030			a| CPTRB4
035			a| LIAS192470
040			d| UtOrBLW
050	0	0	a| TK7885.A1 b| I5
245	0	0	a| IEEE Computer.
247	0	0	f| July 1966-Aug. 1970, a| Computer group news.
264		1	a| Los Angeles, CA..
300			a| volumes : b| illustrations ; c| 28 cm
336			a| text b| txt 2| rdacontent
337			a| unmediated b| n 2| rdamedia
338			a| volume b| nc 2| rdacarrier
362	0		a| Vol. 1 (July 1966)-
550			a| Published by the Institute of Electrical and Electronics Engineers. Computer Society (formerly Computer Group).
650		0	a| Computers v| Periodicals.
710	2		a| IEEE Computer Society.
533			a| Microform. b| [New York, N.Y.? c| IEEE], e| microfiches ; 11 x 15 cm.
949			a| Periodical w| ASIS c| 1 i| 156262-2001 d| 3/2/2001 l| PERIOD-HB m| HARRISBURG r| Y s| Y t| PERIODICAL u| 4/12/2001
949			a| QA76.I56 v.16 no.1-6 1983 w| LCPER c| 1 i| 000018890037 d| 3/2/2001 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 4/12/2001 z| 5054
949			a| QA76.I56 v.16 no.1-6 1983 w| LCPER c| 2 i| 000066946304 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| Microfiche w| ASIS c| 1 i| 156262-9001 d| 3/2/2001 l| MICROS-HB m| HARRISBURG r| Y s| Y t| MICROFORM u| 4/12/2001
949			a| PO-98010864 w| ASIS c| 1 i| PO-98010864 l| ON-ORDER m| SERIAL-SRV r| Y s| Y t| BOOK u| 7/3/2001
949			a| QA76.I56 v.16 no.7-12 1983 w| LCPER c| 1 i| 000052266812 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949			a| QA76.I56 v.16 no.7-12 1983 w| LCPER c| 2 i| 000066946311 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.17 no.1-6 1984 w| LCPER c| 1 i| 000052266829 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949			a| QA76.I56 v.17 no.1-6 1984 w| LCPER c| 2 i| 000052236655 d| 6/1/2004 l| ACAD-BLDG m| UP-ANNEX n| 1 r| Y s| Y t| PERIODICAL u| 5/25/2004
949			a| QA76.I56 v.17 no.7-12 1984 w| LCPER c| 1 i| 000052266836 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949			a| QA76.I56 v.17 no.7-12 1984 w| LCPER c| 2 i| 000066945246 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.19 no.1-6 1986 w| LCPER c| 1 i| 000052266843 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949			a| QA76.I56 v.19 no.1-6 1986 w| LCPER c| 2 i| 000066945277 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t|

ASK A LIBRARIAN

	PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.19 no.7-12 1986 w| LCPER c| 1 i| 000052267031 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.19 no.7-12 1986 w| LCPER c| 2 i| 000066945284 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.20 no.1-6 1987 w| LCPER c| 1 i| 000052267048 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.20 no.1-6 1987 w| LCPER c| 2 i| 000066945291 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.20 no.7-12 1987 w| LCPER c| 1 i| 000052267055 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.20 no.7-12 1987 w| LCPER c| 2 i| 000066945307 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.21 no.1-6 1988 w| LCPER c| 1 i| 000052267062 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.21 no.7-12 1988 w| LCPER c| 1 i| 000052267079 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.22 no.7-12 1989 & suppls. w| LCPER c| 1 i| 000052267093 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.22 no.1-6 1989 w| LCPER c| 1 i| 000052267086 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.23 no.1-6 1990 w| LCPER c| 1 i| 000052267109 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.23 no.7-12 1990 w| LCPER c| 1 i| 000052267116 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.24 no.1-6 1991 w| LCPER c| 1 i| 000052267123 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.24 no.7-12 1991 w| LCPER c| 1 i| 000052267130 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.25 no.1-6 1992 w| LCPER c| 1 i| 000052267147 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.25 no.7-12 1992 w| LCPER c| 1 i| 000052267154 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/23/2003 z| 5054
949	a| QA76.I56 v.26 no.1-6 1993 w| LCPER c| 1 i| 26060101005 d| 5/13/2004 l| CATO-PARK m| UP-ANNEX q| 1 r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 005
949	a| QA76.I56 v.26 no.7-12 1993 w| LCPER c| 1 i| 26060101006 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 006
949	a| QA76.I56 v.27 no.1-6 1994 w| LCPER c| 1 i| 26060101007 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 007
949	a| QA76.I56 v.27 no.7-12 1994 w| LCPER c| 1 i| 26060101008 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 008
949	a| QA76.I56 v.28 no.1-6 1995 w| LCPER c| 1 i| 26060101009 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 009
949	a| QA76.I56 v.28 no.7-12 1995 w| LCPER c| 1 i| 26060101010 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 010
949	a| QA76.I56 v.29 no.1-6 1996 w| LCPER c| 1 i| 26060101011 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 011
949	a| QA76.I56 v.29 no.7-12 1996 w| LCPER c| 1 i| 26060101012 d| 4/7/2010 l| CATO-PARK m| UP-ANNEX q| 1 r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 012
949	a| QA76.I56 v.30 no.1-6 1997 w| LCPER c| 1 i| 26060101013 d| 9/17/2002 l| CATO-PARK m| UP-ANNEX q| 1 r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 013
949	a| QA76.I56 v.30 no.7-12 1997 w| LCPER c| 1 i| 26060101014 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 014
949	a| QA76.I56 v.31 no.1-6 1998 w| LCPER c| 1 i| 26060101015 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 015
949	a| QA76.I56 v.31 no.7-12 1998 w| LCPER c| 1 i| 26060101016 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 016
949	a| QA76.I56 v.32 no.1-6 1999 w| LCPER c| 1 i| 26060101017 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 017
949	a| QA76.I56 v.32 no.7-12 1999 w| LCPER c| 1 i| 26060101018 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 018
949	a| QA76.I56 v.33 no.1-6 2000 w| LCPER c| 1 i| 26060101019 l| CATO-PARK m| UP-ANNEX r| Y s| Y t|

	PERIODICAL u| 5/23/2002 x| D z| 2606 0| 0101 1| 019
949	a| QA76.I56 v.33 no.7-12 2000 w| LCPER c| 1 i| 26060101020 d| 3/2/2001 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 4/12/2001 x| D z| 2606 0| 0101 1| 020
949	a| QA76.I56 v.34 no.1-6 2001 w| LCPER c| 1 i| 26060101021 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 9/26/2002 x| D z| 2606 0| 0101 1| 021
949	a| QA76.I56 v.34 no.7-12 2001 w| LCPER c| 1 i| 26060101022 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 9/26/2002 x| D z| 2606 0| 0101 1| 022
949	a| QA76.I56 v.35 no.1-6 2002 w| LCPER c| 1 i| 26060101023 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/16/2003 x| D z| 2606 0| 0101 1| 023
949	a| QA76.I56 v.35 no.7-12 2002 w| LCPER c| 1 i| 26060101024 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 6/16/2003 x| D z| 2606 0| 0101 1| 024
949	a| QA76.I56 v.36 no.1-6 2003 w| LCPER c| 1 i| 26060101025 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 3/17/2004 x| D z| 2606 0| 0101 1| 025
949	a| QA76.I56 v.36 no.7-12 2003 w| LCPER c| 1 i| 26060101026 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 3/17/2004 x| D z| 2606 0| 0101 1| 026
949	a| QA76.I56 v.37 no.7-12 2004 w| LCPER c| 1 i| 26060101027 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/28/2005 x| D z| 2606 0| 0101 1| 027
949	a| QA76.I56 v.37 no.1-6 2004 w| LCPER c| 1 i| 26060101028 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/28/2005 x| D z| 2606 0| 0101 1| 028
949	a| QA76.I56 v.38 no.1-6 2005 w| LCPER c| 1 i| 26060101029 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/28/2005 x| D z| 2606 0| 0101 1| 029
949	a| QA76.I56 v.38 no.7-12 2005 w| LCPER c| 1 i| 26060101030 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 1/29/2007 x| D z| 2606 0| 0101 1| 030
949	a| QA76.I56 v.39 no.1-6 2006 w| LCPER c| 1 i| 000055099264 l| CATO-2 m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 1/29/2007 z| 6925
949	a| QA76.I56 v.39 no.7-12 2006 w| LCPER c| 1 i| 26060101031 l| CATO-PARK m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 1/29/2007 x| D z| 2606 0| 0101 1| 031
949	a| QA76.I56 v.6 1973 w| LC c| 1 i| 000014186295 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/16/2007
949	a| QA76.I56 v.1 1966/67 w| LCPER c| 1 i| 000014191633 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 8/11/2009
949	a| QA76.I56 v.2 1968/69 w| LCPER c| 1 i| 000066668350 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 8/11/2009
949	a| QA76.I56 v.3 1970 w| LCPER c| 1 i| 000066668367 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 8/11/2009
949	a| QA76.I56 v.4 1971 w| LCPER c| 1 i| 000066668374 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 8/11/2009
949	a| QA76.I56 v.5 1972 w| LCPER c| 1 i| 000066668381 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 8/11/2009
949	a| QA76.I56 v.7 no.1-6 1974 w| LCPER c| 1 i| 000066946120 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.7 no.7-12 1974 w| LCPER c| 1 i| 000066946137 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.8 no.1-6 1975 w| LCPER c| 1 i| 000066946144 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.8 no.7-12 1975 w| LCPER c| 1 i| 000066946151 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.9 no.1-6 1976 w| LCPER c| 1 i| 000066946168 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.9 no.7-12 1976 w| LCPER c| 1 i| 000066946175 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.10 no.1-6 1977 w| LCPER c| 1 i| 000066946182 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.10 no.7-12 1977 w| LCPER c| 1 i| 000066946199 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.11 no.1-6 1978 w| LCPER c| 1 i| 000066946205 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.11 no.7-12 1978 w| LCPER c| 1 i| 000066946212 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.12 no.1-6 1979 w| LCPER c| 1 i| 000066946229 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949	a| QA76.I56 v.12 no.7-12 1979 w| LCPER c| 1 i| 000066946236 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t|

			PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.13 no.1-6 1980 w| LCPER c| 1 i| 000066946243 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.13 no.7-12 1980 w| LCPER c| 1 i| 000066946250 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.14 no.1-6 1981 w| LCPER c| 1 i| 000066946267 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.14 no.7-12 1981 w| LCPER c| 1 i| 000066946274 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.15 no.1-6 1982 w| LCPER c| 1 i| 000066946281 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.15 no.7-12 1982 w| LCPER c| 1 i| 000066946298 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.18 no.1-6 1985 w| LCPER c| 1 i| 000066945253 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
949			a| QA76.I56 v.18 no.7-12 1985 w| LCPER c| 1 i| 000066945260 l| ACAD-BLDG m| UP-ANNEX r| Y s| Y t| PERIODICAL u| 10/6/2009
852	0	1	a| PSt b| WILKESBAR c| STACKS-WB h| Periodical
866		0 8	0 a| Current Year Only.
852	0	1	a| PSt b| UP-ANNEX c| ANNEX h| QA76 i| .I56
866		0 8	0 a| v.1(1966)-v.39(2006)
852	8		a| PSt b| HARRISBURG c| MICROS-HB h| Microfiche
866		0 8	0 a| v.13 1979-v.32 1999
852	8		a| PSt b| HARRISBURG c| PERIOD-HB h| Periodical
866		0 8	0 a| v.5 1972-v.12 1979

PennState

Copyright @2019 The Pennsylvania State University. All rights reserved. Except where otherwise noted, this work is subject to a Creative Commons Attribution 4.0 license. Details and exceptions.

Legal Statements | PSU Hotlines

PENN STATE UNIVERSITY LIBRARIES

Libraries Home (814) 865-6368
Libraries Intranet (Staff Only)
Accessibility Help
Website Feedback
Policies and Guidelines
Staff Directory

CONNECT WITH PENN STATE UNIVERSITY LIBRARIES

Facebook
Twitter
Instagram

PennState University Libraries | Catalog

Bookmarks 0 Course Reserves My Account

Keyword Search... Search Advanced search Start Over

Share Bookmark Report an Issue

IEEE Computer

Additional Titles	July 1966-Aug. 1970, Computer group news
Published	Los Angeles, CA..
Physical Description	volumes : illustrations ; 28 cm
Additional Creators	IEEE Computer Society

ASK A LIBRARIAN

Availability

Browse Nearby on Shelf I Want It

Penn State Harrisburg (2 items)

Call number	Material	Location
Harrisburg - Microforms - 2nd Floor - In-Library Use: Holdings Summary v.13 1979-v.32 1999 Harrisburg - Periodicals - 2nd Floor - In-Library Use: Holdings Summary v.5 1972-v.12 1979		
Microfiche	Microfilm, Microfiche, etc.	Harrisburg - Microforms - 2nd Floor - In-Library Use
Periodical	Bound Journal	Harrisburg - Periodicals - 2nd Floor - In-Library Use

Library Storage (80 items)

Call number	Material	Location
Annexed Material: Holdings Summary v.1(1966)-v.39(2006)		
QA76.I56 v.1 1966/67	Bound Journal	Annexed Material
QA76.I56 v.2 1968/69	Bound Journal	Annexed Material
QA76.I56 v.3 1970	Bound Journal	Annexed Material
QA76.I56 v.4 1971	Bound Journal	Annexed Material
QA76.I56 v.5 1972	Bound Journal	Annexed Material
QA76.I56 v.6 1973	Bound Journal	Annexed Material
QA76.I56 v.7 no.1-6 1974	Bound Journal	Annexed Material
QA76.I56 v.7 no.7-12 1974	Bound Journal	Annexed Material
QA76.I56 v.8 no.1-6 1975	Bound Journal	Annexed Material
QA76.I56 v.8 no.7-12 1975	Bound Journal	Annexed Material
QA76.I56 v.9 no.1-6 1976	Bound Journal	Annexed Material
QA76.I56 v.9 no.7-12 1976	Bound Journal	Annexed Material
QA76.I56 v.10 no.1-6 1977	Bound Journal	Annexed Material
QA76.I56 v.10 no.7-12 1977	Bound Journal	Annexed Material
QA76.I56 v.11 no.1-6 1978	Bound Journal	Annexed Material
QA76.I56 v.11 no.7-12 1978	Bound Journal	Annexed Material

QA76.I56 v.12 no.1-6 1979	Bound Journal	Annexed Material
QA76.I56 v.12 no.7-12 1979	Bound Journal	Annexed Material
QA76.I56 v.13 no.1-6 1980	Bound Journal	Annexed Material
QA76.I56 v.13 no.7-12 1980	Bound Journal	Annexed Material
QA76.I56 v.14 no.1-6 1981	Bound Journal	Annexed Material
QA76.I56 v.14 no.7-12 1981	Bound Journal	Annexed Material
QA76.I56 v.15 no.1-6 1982	Bound Journal	Annexed Material
QA76.I56 v.15 no.7-12 1982	Bound Journal	Annexed Material
QA76.I56 v.16 no.1-6 1983	Bound Journal	Annexed Material
QA76.I56 v.16 no.1-6 1983	Bound Journal	Annexed Material
QA76.I56 v.16 no.7-12 1983	Bound Journal	Annexed Material
QA76.I56 v.16 no.7-12 1983	Bound Journal	Annexed Material
QA76.I56 v.17 no.1-6 1984	Bound Journal	Annexed Material
QA76.I56 v.17 no.1-6 1984	Bound Journal	Annexed Material
QA76.I56 v.17 no.7-12 1984	Bound Journal	Annexed Material
QA76.I56 v.17 no.7-12 1984	Bound Journal	Annexed Material
QA76.I56 v.18 no.1-6 1985	Bound Journal	Annexed Material
QA76.I56 v.18 no.7-12 1985	Bound Journal	Annexed Material
QA76.I56 v.19 no.1-6 1986	Bound Journal	Annexed Material
QA76.I56 v.19 no.1-6 1986	Bound Journal	Annexed Material
QA76.I56 v.19 no.7-12 1986	Bound Journal	Annexed Material
QA76.I56 v.19 no.7-12 1986	Bound Journal	Annexed Material
QA76.I56 v.20 no.1-6 1987	Bound Journal	Annexed Material
QA76.I56 v.20 no.1-6 1987	Bound Journal	Annexed Material
QA76.I56 v.20 no.7-12 1987	Bound Journal	Annexed Material
QA76.I56 v.20 no.7-12 1987	Bound Journal	Annexed Material
QA76.I56 v.21 no.1-6 1988	Bound Journal	Annexed Material
QA76.I56 v.21 no.7-12 1988	Bound Journal	Annexed Material
QA76.I56 v.22 no.1-6 1989	Bound Journal	Annexed Material
QA76.I56 v.22 no.7-12 1989 & suppls.	Bound Journal	Annexed Material
QA76.I56 v.23 no.1-6 1990	Bound Journal	Annexed Material
QA76.I56 v.23 no.7-12 1990	Bound Journal	Annexed Material
QA76.I56 v.24 no.1-6 1991	Bound Journal	Annexed Material
QA76.I56 v.24 no.7-12 1991	Bound Journal	Annexed Material
QA76.I56 v.25 no.1-6 1992	Bound Journal	Annexed Material
QA76.I56 v.25 no.7-12 1992	Bound Journal	Annexed Material
QA76.I56 v.26 no.1-6 1993	Bound Journal	Annexed Material
QA76.I56 v.26 no.7-12 1993	Bound Journal	Annexed Material
QA76.I56 v.27 no.1-6 1994	Bound Journal	Annexed Material
QA76.I56 v.27 no.7-12 1994	Bound Journal	Annexed Material
QA76.I56 v.28 no.1-6 1995	Bound Journal	Annexed Material
QA76.I56 v.28 no.7-12 1995	Bound Journal	Annexed Material
QA76.I56 v.29 no.1-6 1996	Bound Journal	Annexed Material
QA76.I56 v.29 no.7-12 1996	Bound Journal	Annexed Material

QA76.I56 v.30 no.1-6 1997	Bound Journal	Annexed Material
QA76.I56 v.30 no.7-12 1997	Bound Journal	Annexed Material
QA76.I56 v.31 no.1-6 1998	Bound Journal	Annexed Material
QA76.I56 v.31 no.7-12 1998	Bound Journal	Annexed Material
QA76.I56 v.32 no.1-6 1999	Bound Journal	Checked out
QA76.I56 v.32 no.7-12 1999	Bound Journal	Checked out
QA76.I56 v.33 no.1-6 2000	Bound Journal	Annexed Material
QA76.I56 v.33 no.7-12 2000	Bound Journal	Annexed Material
QA76.I56 v.34 no.1-6 2001	Bound Journal	Annexed Material
QA76.I56 v.34 no.7-12 2001	Bound Journal	Annexed Material
QA76.I56 v.35 no.1-6 2002	Bound Journal	Annexed Material
QA76.I56 v.35 no.7-12 2002	Bound Journal	Annexed Material
QA76.I56 v.36 no.1-6 2003	Bound Journal	Annexed Material
QA76.I56 v.36 no.7-12 2003	Bound Journal	Annexed Material
QA76.I56 v.37 no.1-6 2004	Bound Journal	Annexed Material
QA76.I56 v.37 no.7-12 2004	Bound Journal	Annexed Material
QA76.I56 v.38 no.1-6 2005	Bound Journal	Annexed Material
QA76.I56 v.38 no.7-12 2005	Bound Journal	Annexed Material
QA76.I56 v.39 no.1-6 2006	Bound Journal	Annexed Material
QA76.I56 v.39 no.7-12 2006	Bound Journal	Annexed Material

Acquisitions (UP) (1 item)

Call number	Material	Location
PO-98010864	Book	Being acquired by the library

Dates of Publication and/or Sequential Designation
Vol. 1 (July 1966)-

Subject(s)
Computers—Periodicals

ISSN
0018-9162

Reproduction Note
Microform. [New York, N.Y.? IEEE], microfiches ; 11 x 15 cm.

Issuing Body
Published by the Institute of Electrical and Electronics Engineers. Computer Society (formerly Computer Group).

View MARC record | catkey: 156262

PennState

Copyright @2019 The Pennsylvania State University. All rights reserved. Except where otherwise noted, this work is subject to a Creative Commons Attribution 4.0 license. Details and exceptions.

Legal Statements | PSU Hotlines

PENN STATE UNIVERSITY LIBRARIES
Libraries Home (814) 865-6368
Libraries Intranet (Staff Only)
Accessibility Help
Website Feedback
Policies and Guidelines
Staff Directory

CONNECT WITH PENN STATE UNIVERSITY LIBRARIES
Facebook
Twitter
Instagram

APPENDIX LAWTON02


PSU LIBRARIES ANNEX
2606 01 01 017


COMPUTER
32
JAN-JUNE
1999
ENGIN
QA76
.I56

A0000478786D0

ENGINEERING LIBRARY

SHELVED IN
BOUND
PERIODICAL
COLLECTION

COMPUTER

Innovative Technology for Computer Professionals

March 1999



The Crypto Community: Strength in Numbers

Taking the Network Home

Announcing

IT Professional

TECHNOLOGY SOLUTIONS FOR THE ENTERPRISE

see page 7

********** 5-DIGIT 16802
031999 300 00014
PENN STATE UNIV
PATTEE LIB SER RECORDS
UNIVERSITY PARK PA 16802

IEEE COMPUTER SOCIETY

http://computer.org

You Make the Call: Ethics in the Market, p. 80

March 1999, Volume 32, Number Three

IEEE Computer Society: http://computer.org
Computer: http://computer.org/computer
computer@computer.org

COMPUTER
innovative technology for computer professionals

NEWS

11 Industry Trends
New Technologies Take the Network Home
George Lawton

15 Technology News
Will 1999 Be the Year of IP Telephony?
Neal Leavitt

18 News Briefs
Groups Duel over New I/O Standards ❖ Oracle, Sun Team Up to Dump the OS ❖ Bringing the Net to Mobile Appliances



A growing number of vendors are selling a variety of home-networking products.

Y2K REPORT

21 Y2K in Brief
US Corporations' Y2K Efforts Go Global

22 Rubin's Y2K Diary
Doing the Y2K Spin: Is Image Everything?

DEPARTMENTS

4 Letters

103 New Books

104 New Tools



Has R&D's dependence on deep-pocket funding sacrificed innovation for quick payoffs?

COLUMNS

108 Internet Watch
Cryptography: The Importance of Not Being Different
Bruce Schneier

110 Standards
Defining Stakeholder Relationships
Harold W. Lawson

113 Management
Making RAD Work for Your Project
Barry Boehm

115 Component and Object Technology
Scientific Components Are Coming
Paul F. Dubois

120 Binary Critic
Innovation in the Small
Ted Lewis

MEMBERSHIP NEWS

76 Computer Society Connection
Society Activities Attract New Members ❖ Two Presented Technical Achievement Awards ❖ Michael C. Mulder Tribute

80 Your Call
When a Product Doesn't Measure Up to Its Claims

83 Call and Calendar

Computer's NEW TOOLS

http://computer.org/computer/tools

IEEE CS Membership Application, 8
Advertiser/Product Index, 101
Career Opportunities, 92

24 **Article Summaries**



CYBERSQUARE

27 **Have We Witnessed a Real-Life Turing Test?**

Marina Krol

COMPUTING PRACTICES

31 **Beyond Spreadsheets: Tools for Building Decision Support Systems**

Hemant K. Bhargava, Suresh Sridhar, and Craig Herrick

COVER FEATURES



40 **Transforming Business Through Information Technology**

Girish Pathak and Shri K. Goyal

42 **The Changing Role of Information Technology in Manufacturing**

Krishnamurthy Srinivasan and Sundaresan Jayaraman

50 **Reinventing GTE with Information Technology**

Richard Brandau, Tony Confrey, Alin D'Silva, Christopher Matheus, and Robert Weihmayer

59 **The BT Intranet: Information by Design**

Mark Norris, Paul Muschamp, and Steve Sim

67 **Making Business Sense of Electronic Commerce**

Dawn Jutla, Peter Bodorik, Catherine Hajnal, and Charles Davis

COPYRIGHT © 1999 BY THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS INC. ALL RIGHTS RESERVED. ABSTRACTING IS PERMITTED WITH CREDIT TO THE SOURCE. LIBRARIES ARE PERMITTED TO PHOTOCOPY BEYOND THE LIMITS OF US COPYRIGHT LAW FOR PRIVATE USE OF PATRONS: (1) THOSE POST-1977 ARTICLES THAT CARRY A CODE AT THE BOTTOM OF THE FIRST PAGE, PROVIDED THE PER-COPY FEE INDICATED IN THE CODE IS PAID THROUGH THE COPYRIGHT CLEARANCE CENTER, 222 ROSEWOOD DR., DANVERS, MA 01923; (2) PRE-1978 ARTICLES WITHOUT FEE. FOR OTHER COPYING, REPRINT, OR REPUBLICATION PERMISSION, WRITE TO COPYRIGHTS AND PERMISSIONS DEPARTMENT, IEEE PUBLICATIONS ADMINISTRATION, 445 HOES LANE, P.O. BOX 1331, PISCATAWAY, NJ 08855-1331.

EDITOR-IN-CHIEF:
JAMES H. AYLOR, UNIV. OF VIRGINIA, JHA@VIRGINIA.EDU (804) 924-6100

ASSOCIATE EDITOR-IN-CHIEF:
DORIS CARVER, LOUISIANA STATE UNIV., CARVER@BIT.CSC.LSU.EDU

COMPUTING PRACTICES:
THOMAS CAIN, UNIV. OF PITTSBURGH; T.CAIN@COMPUTER.ORG

CYBERSQUARE:
RONALD HOELZEMAN, UNIV. OF PITTSBURGH, R.HOELZEMAN@ COMPUTER.ORG

SPECIAL ISSUES:
KATHLEEN SWIGGER, UNIV. OF NORTH TEXAS, KATHY@CS.UNT.EDU

1999 PRESIDENT:
LEONARD TRIPP

EDITORIAL STAFF

ANGELA BURGESS
EXECUTIVE EDITOR
A.BURGESS@COMPUTER.ORG

SCOTT HAMILTON
SENIOR ACQUISITIONS EDITOR
S.HAMILTON@COMPUTER.ORG

JANET WILSON
SENIOR ACQUISITIONS EDITOR
J.WILSON@COMPUTER.ORG

LEE GARBER
NEWS EDITOR
L.GARBER@COMPUTER.ORG

CHRISTINE MILLER
PRODUCTION EDITOR
C.MILLER@COMPUTER.ORG

JAMES SANDERS
ASSOCIATE EDITOR

KIRK L. KROEKER
STAFF EDITOR

MARY-LOUISE G. PINER
STAFF EDITOR, MEMBERSHIP NEWS

GREG GOTH
STAFF EDITOR, NEWS

JASON SEABORN
EDITORIAL ASSISTANT

MICHELLE J. SAEWERT
ELECTRONIC MEDIA ASSISTANT

SHANNON BUCKELS
PRODUCTION ASSISTANT

HEATHER SEMON
MAGAZINE ASSISTANT

EXECUTIVE DIRECTOR & CHIEF EXECUTIVE OFFICER
T. MICHAEL ELLIOTT

PUBLISHER
MATT LOEB
M.LOEB@COMPUTER.ORG

MEMBERSHIP & CIRCULATION MARKETING MANAGER
GEORGANN CARTER

ADVERTISING MANAGER
PATRICIA GARVEY

ADVERTISING SUPERVISOR
MARIAN ANDERSON

DESIGN
DIRK HAGNER

PRODUCTION
CHRISTOPHER A. PATTERSON
LARRY BAUER
KEN DUCKWORTH

NEXT MONTH: Active Networks

MEMBERSHIP MAGAZINE OF THE IEEE COMPUTER SOCIETY

Cover design and artwork by Dirk Hagner

BPA INTERNATIONAL

New Technologies Take the Network Home

George Lawton

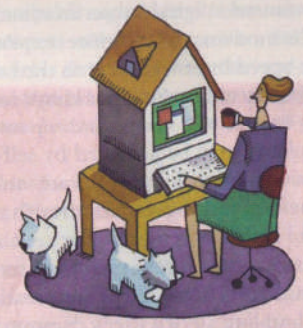
Home networks, once the realm of the rich and technically gifted, are now considered to be one of the computer industry's hottest markets. Although home networking is still in its early stages, a number of trends have made this technology important.

During the past couple of years, an increasing number of households have acquired more than one PC. To permit the users of different computers in a home to access the same files, play video games against each other, and share Internet access and peripherals, a growing number of consumers are networking their homes.

And the increasing use of embedded chips in domestic appliances (as well as in security, temperature control, and other household systems) that can connect to a home network has accelerated this process.

In response to this trend and the availability of a potentially lucrative untapped market, companies have developed technologies for low-cost, easy-to-use home networks.

Until recently, home networking was restricted to low-speed technologies or traditional LANs, which are too difficult for most consumers to install and maintain. However, this is changing because of new, faster, easier-to-use home-networking approaches, enabled in part by more powerful and less expensive processor technology.



A growing number of vendors are selling a variety of home-networking products.

Major vendors in the computer and consumer-electronics industries are selling home-networking products, including chips (such as stand-alone and Ethernet chips) and networking cards (such as PCI cards), that use phone, wireless, coaxial cable, and power-line technology.

Meanwhile, consortiums and companies like Microsoft and Sun Microsystems are creating architectures that will let PCs and appliances on different platforms communicate with each other over home networks.

Because of these factors, said analyst Kurt Scherf with Parks Associates, a market research firm, revenues from the sale of home-networking products could grow from $10 million in 1998 to almost $200 million in 2001 to about $375 million in 2004, as Figure 1 shows.

Home networks currently reach speeds of up to 10 megabits per second (Mbps) and thus are slower than corporate networks. This is the case primarily because home networks are built over phone lines and other media that are easier to maintain but of lower quality than those used in traditional LANs. However, home-networking vendors say they are working on products that will transmit at rates of up to 100 Mbps.

Meanwhile, home-networking technologies should be able to link devices in all but the largest homes, since even the technology with the shortest range still works at distances of 150 feet.

Home-networking analyst Karuna Uppal with the Yankee Group, a market research firm, said users will begin to adopt the technology in earnest this year.

Parks Associates estimates the total number of households with LANs will grow from 2.5 million in 1998 to 4.3 million next year to 11.9 million by 2004.

Meanwhile, the technology will also be attractive to smaller businesses, particularly those without IT departments and those in older buildings without fiber or other high-quality wiring.

BACKGROUND

People have been talking about home networking for more than a decade. Initially, home networking was a slow technology used just for turning on and off lights and household appliances, as well as other home automation tasks.

However, lack of enthusiasm by consumers inhibited the technology, limited product development, and left home networking primarily to users with a lot of money or technical talent.

Now, however, less expensive and more powerful CPUs are permitting PCs to act like servers for home networks and like gateways for distributing Internet access. This is the primary reason that the

Editor: Lee Garber, *Computer*, 10662 Los Vaqueros Circle, PO Box 3014, Los Alamitos, CA 90720-1314; l.garber@computer.org

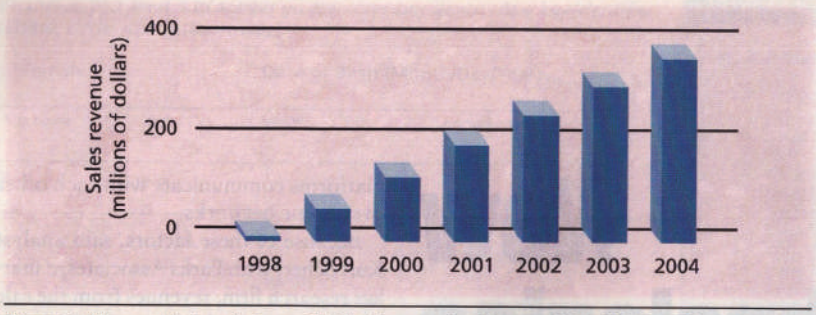


Figure 1. Many market analysts predict the home-networking market will take off in 1999. Parks Associates estimates revenue from the sale of home-networking products will soar from about $10 million last year to about $75 million this year to about $200 million in 2001.

home networking industry is growing now, noted Craig Stouffer, vice president of marketing at networking vendor Tut Systems.

GROWING DEMAND

The demand for home networking is being driven by the growth of multicomputer homes, which is being driven by the growing number of computer users and low-cost computers.

Dataquest, a market research firm, said the number of US homes with multiple PCs will grow from 12.3 million in 1997 to 20.1 million in 2001, as Figure 2 shows. However, noted analyst Bruce Ryan with Media Metrix, a computer-industry market-measurement firm, some multicomputer households don't use more than one of their PCs.

The Yankee Group's Uppal said

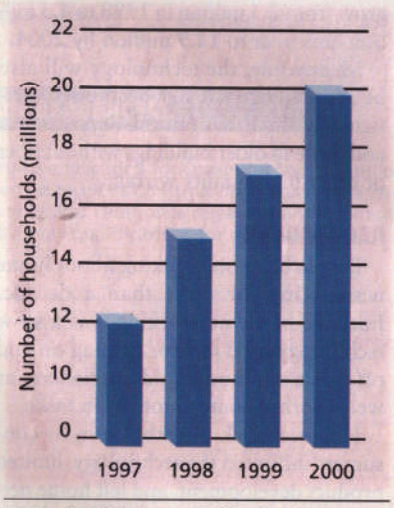


Figure 2. Statistics and estimates from Dataquest say the number of US homes with more than one PC has increased steadily during the past couple of years and will continue to do so. Many industry observers say that the growing number of multi-PC homes is driving the demand for home-networking products.

Internet access has been the primary source of interest in home networks. Home networks can transmit Internet-based video, audio, and other multimedia presentations from a gateway computer to other PCs, TVs, and entertainment centers. It is less expensive for users to let these devices share a single Internet account than to set up multiple accounts.

These trends will increase as ADSL (asymmetric digital subscriber line) and cable modems bring more expensive, high-speed Internet access to the home.

The key to home networking's success will be simple setup and operation. Simplicity may be achieved by self-configuring networking software and by home-networking capabilities that are built into computers. For example, Compaq sells computers with its Net Wizard software that can automatically find and link all of a user's computers.

HOME-NETWORKING TECHNOLOGIES

Many industry observers predict that most home-networking technologies will operate over phone lines. However, vendors are also working with wireless, cable, and power-line technologies.

Home-networking technologies work with architectures that let applications find and configure devices automatically, which eliminates the setup and configuration necessary with corporate networks.

Phone lines

Phone-line networks appear particularly promising because telephone service is so widespread and because data transmissions over phone wiring experience relatively little interference from analog voice transmissions. Frequency-division multiplexing divides a phone line into separate communications channels, each with a different frequency. This permits data, Internet, and voice communications to use the same line simultaneously. With power-line or wireless technologies, on the other hand, there can be considerable interference.

Meanwhile, phone-line networking is less expensive than other home-networking technologies. In addition, said Uppal, phone-line networking promises higher speeds in the near future.

No hubs. Unlike traditional corporate LANs that operate over telephone wiring, home phone-line networks do not have to be set up in a star topology, as Figure 3 shows, because participating PCs and other devices can recognize each other without a hub. Rather, users can set up networks by simply plugging devices with network cards into telephone outlets. Thus, phone-line networking saves consumers from having to buy hubs or configure their networks.

HomePNA. The Home Phoneline Networking Association (http://www.homepna.org) includes such companies as Compaq, Intel, and Rockwell. The HomePNA is working on standards to ensure compatibility between home-networking products and has developed a Home Phoneline Network seal of approval for companies that comply, said Jeff Thermond, CEO of Epigram, a home-networking company that is a HomePNA member.

The consortium is focusing on similar technologies developed by Epigram and Tut Systems. The HomePNA has released its first specification based on Tut's 1-Mbps Ethernet-based HomeRun technology. Such companies as AMD, Intel, Lucent Technologies, National Semiconductor, and Rockwell are manufacturing HomeRun-based chips to be used in modems and networking cards.

The HomePNA is working with Epigram to develop a 10-Mbps specification. Epigram already has its iLine10 10-Mbps Ethernet LAN chipset technology, which is backward compatible with the 1-Mbps technology used in the current HomePNA specification. 3Com and Nortel Networks' Netgear subsidiary have already released iLine10-based products, including network cards.

Meanwhile, Avio Digital is working on an 88-Mbps technology, while Epigram is working on a 100-Mbps technology.

Limitations. One potential limitation of

phone-line-based home networks is that they can support a networked device only where there is a phone jack. This is not necessarily a critical limitation in the US, where, according to Intel consumer research, the average home has 4.5 phone jacks. However, homes in most other countries do not have as many jacks.

Wireless

A number of companies—such as Diamond Multimedia Systems, Proxim, Sharewave, and Web Gear—are selling wireless home-networking products. Wireless home networks provide users with convenience by opening up communications capabilities to mobile devices.

Wireless networks use receivers and transmitters that operate over radio or infrared waves. The networks transmit at rates from 1 to 4 Mbps. Infrared-based products work over several dozen feet, while radio-based products work over distances up to several hundred feet.

Most wireless systems tend to be more expensive than other home-networking systems because they must be engineered with relatively powerful radio transmitters and receivers, as well as sophisticated digital signal processors. DSPs filter out some of the interference that comes from other devices, such as phones, and that can slow or even shut down a network.

Convergence. Some companies are looking to wireless networking to promote the convergence of home computing and home entertainment. For example, Philips Electronics' Philips Peripherals Group is developing Ambi. Ambi will use Sharewave's wireless-networking technology in set-top boxes that will let consumers access computer functionality from their televisions, using a kind of TV browser.

Interoperability. A current limitation of wireless networking is that different vendors' products are not interoperable. Proxim has launched the Wireless LAN Interoperability Forum (http://www.wlif.org) and is working with 25 other companies, including IBM and Motorola, to set standards for this technology. The standards would be based on Proxim's radio-frequency (RF) LAN technology, which permits transmission speeds of up to 1.6 Mbps.

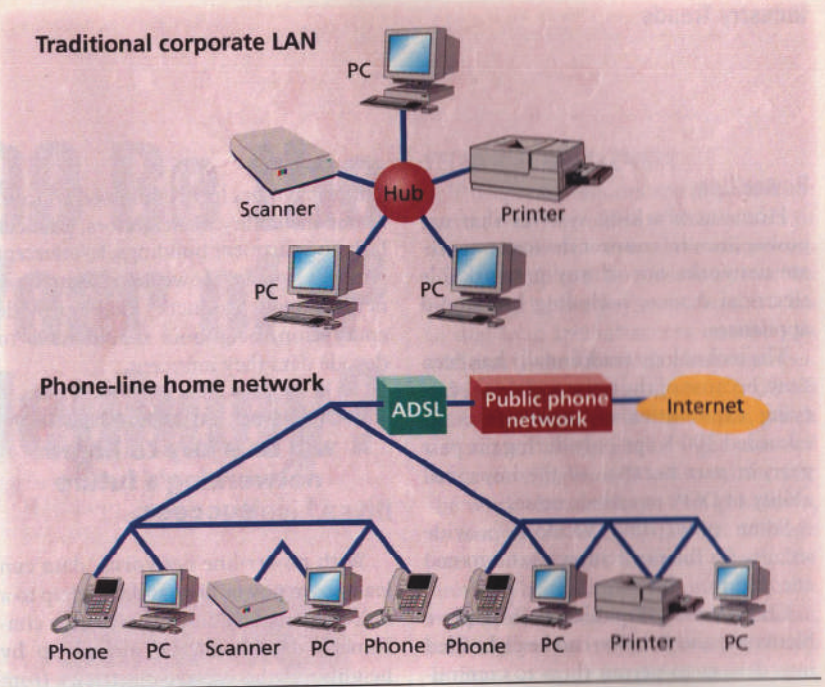


Figure 3. In a traditional corporate star-topology LAN, each device is linked to a port on a hub. This requires organizations to buy hubs and configure networks. In a phone-line home network, all devices with network cards are linked via phone lines (as long as they are within a maximum allowable distance of each other) and automatically recognize each other. This eliminates the need for a hub and makes network operation simpler and less expensive. Home networks can connect via high-speed ADSL modems to the Internet via the public phone network.

Meanwhile, such companies as Compaq, Hewlett-Packard, Intel, Microsoft, and Sony are working with the HomeRF Working Group (http://www.homerf.org) to develop the Shared Wireless Access Protocol (SWAP) for radio-based home networks. The specification allows data rates up to 2 Mbps, lets home networks accommodate up to 127 nodes, and transmits up to 150 feet through walls.

HomeRF Working Group Chair Ben Manny said the first SWAP-based products to emerge later this year will probably permit the operation of networks that let laptops share resources and Internet connectivity with other computers.

He said the working group has also been looking at a lower strength SWAP Lite that could be used for lower cost connectivity applications.

Cable

The coaxial cable installed in many homes for use with TVs can also carry data signals and, in fact, can carry a signal of more bandwidth over a longer distance with less attenuation than other home-networking media.

However, said Hongjun Li, Parks Associates' director of research, only 67 percent of US households have cable TV service, while 96 percent have telephone service. In addition, consumers have little experience and expertise adding new cable for purposes other than receiving TV signals. Therefore, relatively few companies are pursuing cable-based home networking.

In most cable-based home networks, set-top boxes serve as gateways for connecting devices to each other or to the Internet. The boxes would connect to other devices via, for example, infrared transmissions or the IEEE 1394 high-performance serial bus.

Consumers could use the boxes to tune in video and audio channels on the cable TV network or the Internet. The hubs would translate the signals into a form that could be distributed over the home network to the various receiving devices. Cable modems could be used for high-speed Internet access.

Cable TV set-top-box vendor General Instrument is working with Sony's Aperios networking software to deliver home-networking services. Another vendor, Scientific Atlanta, is working to integrate Avio Digital's proposed 88-Mbps MediaWire home-networking technology into its Explorer 2000 set-top box.

Power line

Home-networking systems that use power lines to connect devices can create networks out of any networkable electrical device, including household appliances.

The technology traditionally has been slow, because of the noise and signal-carrying capabilities of power lines. It exceeded 300 Kbps only during the past year, in part because of the improved ability of DSPs to reduce noise.

Some power-line vendors provide security for their systems with end-to-end encryption.

Chips from companies like Adaptive Networks and Echelon can be embedded into devices to permit them to communicate directly over power lines at lower speeds.

Intelogis recently released the Passport Plug-In Network. The system uses adapters, which are external networking modules that plug into electrical outlets and then connect to networked devices via parallel cables. The product currently supports a 350-Kbps network that links devices up to a quarter-mile apart. Intelogis plans to have a 2-Mbps version available by the end of this year.

Todd Green, Intelogis' director of marketing, said power-line networking's main advantages are that homes generally have more electrical outlets than phone outlets and that, unlike wireless systems, power lines allow transmission through concrete or metal floors.

However, power-line technology has been limited by such technical issues as potentially high interference levels from electrical transmissions. In addition, homes can have various configurations of splitters and/or outlets, which means their electrical properties can differ.

OBSTACLES, CHALLENGES, AND PROBLEMS

Although apparently a rising star in the computer industry, home networking still faces its share of challenges to widespread adoption.

Security

Park Associates' Li said users appear to have confidence in the security of wired home-networking technology. Li said some users have expressed concern about the ability of neighbors, particularly in apartment buildings, to intercept wireless signals. However, if users have networking products that provide encryption, neighbors would have to decode data they intercept.

Increased interoperability will be a key to home networking's future success.

With power-line networks, data can leak down power lines and travel up to a half-mile from a house. Therefore, a consumer's data could be picked up by neighbors who receive electricity from the same power line and use the same networking equipment. Again, though, neighbors would have to decode data that has been encrypted.

Interoperability

Several efforts are under way to tie home-networking applications together across platforms and media.

HomeAPI. The HomeAPI Working Group is developing an open industry specification that defines a standard set of software services and application programming interfaces that let applications discover, monitor, and control networked devices and systems.

The working group expects to release the first HomeAPI developers kit for Windows this year and to support other operating systems later.

HAVi. Eight Japanese and European consumer-electronics vendors—including Germany's Grundig, France's Thomson Multimedia, and Japan's Toshiba—have developed the HAVi (home audio/video interoperability) specification primarily for connecting digital home-entertainment devices, although it can also be used with PCs.

The specification lets devices communicate via their IEEE-1394-compatible high-performance serial buses. A number of industry observers have expressed interest in the standard but say it should work with all devices, not just those with IEEE 1394 buses.

Avio Digital CEO Eugene Van Bergen said, "I think it is always undesirable to link an architectural standard like that to a particular media."

Universal Plug-and-Play. Microsoft's Universal Plug-and-Play initiative is designed to let users connect PCs, networked appliances, networks, services, and so on without complicated setup or configuration. In addition, a PC would not be needed for other devices to connect.

Universal Plug-and-Play would define a set of common interfaces that describe devices and services, and permit automatic discovery when a device connects to a network. It would not be restricted to the Windows platform, said Alec Saunders, a Microsoft group planning manager.

Jini. Sun's Jini technology would run a set of Java class libraries on top of various enabled devices' Java virtual machines (JVMs) and would also specify APIs for various hardware infrastructures and operating systems. This would permit enabled devices using various platforms to form a federation of JVMs that could communicate with each other.

Parks Associates' Scherf said household LAN products will begin to proliferate this year and will then rapidly begin to reach a mass market.

A key factor will be the success of the interoperability initiatives. When home-networking products are interoperable, consumers can buy the networking components that suit their needs and pocketbooks. This, in turn, will increase competition, which could lead to a bigger variety of better, less expensive products.

Meanwhile, said Bob Dillon, Enikia's vice president of marketing, the next generation of broadband network services may encourage even more users to buy home-networking systems.

Said Scherf, "The rapid growth of multi-PC households, the proliferation of the Internet, and the complications of traditional computer networking solutions have set the perfect stage for the birth of the home-networking industry." ❖

George Lawton is a freelance technology writer based in Brisbane, California. Contact him at glawton@best.com.

ENGINEERING LIBRARY

