

**Design and Implementation of an Ubiquitous High Speed Communication Utility for
Carnegie-Mellon University**

John Leong

Information Technology Center - Computing Center

Carnegie-Mellon University

Schenley Park, Pittsburgh, PA 15213

412 268 6722

April, 1986

Abstract

CMU is a small private university with roughly 5,500 students and 1,500 faculty and staff. Computing has played a significant role for all departments - from the traditional multi-discipline academic and research computing to administrative support functions. Within the past 5 years, the computing environment at CMU has been undergoing major evolution from the central data center model to one of distributed processing. In 1982, CMU's commitment to distributed computing was formalised with the commencement of a joint project with IBM. An important aspect of the project is to provide the campus with a comprehensive and saturating network utility of more than 10,000 high speed communication ports. This paper describes how we are migrating from the existing situation towards the provision of a ubiquitous, high performance network utility which is resilient to technological changes for the campus. It addresses the issues of local area networking technologies, communication protocols, LAN interconnections and cable plant selection.

Introduction

CMU is a small private university with roughly 5,500 students and 1,500 faculty and staff. Computing has played a significant role for all departments - from the traditional multi-discipline academic and research computing to administrative support functions such as document preparation, spreadsheet and database applications. Electronic mail and bulletin boards are also widely used. Within the past 5 years, the computing environment at CMU has been undergoing major evolution from the central data center model to one of distributed processing. Until the late 70's, with few exceptions such as the Computer Science department, most departments relied on the computing resources provided by the Computation Center. Since then, with the decreasing cost of computing hardware, individual departments started acquiring their own machines. The pace of the computing power distribution accelerated with the introduction of high powered work stations and personal computers. Currently, there are over 150 DEC VAX 750/780 computers on campus together with approximately 200 SUN's, 100 RT-PC's, 100 microVAX'es, 200 ICL Perq's, 50 DEC Pro 350's, 50 HP9836's, numerous DEC PDP-11's and miscellaneous work stations such as Xerox Altos, and Symbolics's LISP machines. In addition to the above, there are some 1500 IBM PC's and 500 Apple Macintoshes around the campus. The Computing Center operates 6 DEC-20's, 5 VAX's and a recently acquired an IBM 3083.

In 1982, CMU's commitment to distributed computing was formalised with the commencement of a joint project with IBM. An important aspect of the project is to provide the campus with a comprehensive and saturating network utility of more than 10,000 high speed communication ports as well as the development of a large scale distributed filing system.

Prior to this project, CMU already had a significant amount of data communication capabilities. There is a fibre optic backbone network consisting of 150 50 micron fibre cables configured in a star topology reaching most of the major buildings on campus. This fibre optic plant was used primarily to support a large multiplexed 2000 line and 1000 port terminal network controlled by 3 Micom data switches. In addition, there are a number of high speed local area networks (Ethernets and ProNets) operated by different departments across the campus.

This paper describes how we are migrating from the existing situation towards the

provision of a ubiquitous, high performance network utility which is resilient to technological changes for the campus. It addresses the issues of local area networking technologies, communication protocols, LAN interconnections and cable plant selection.

Why local area network

Traditionally, most communication traffic is explicitly generated by users. The most common application is host access from terminals, followed by file transfer and mail. These types of applications can normally be satisfied to a large extent by the standard serial line network especially if it operates at the relatively high speed of 9.6K or 19.2K bps. Even when a user requests the transfer of a large file, the delay is generally tolerable since the transfer is explicitly requested and the user is psychologically primed to wait. For those infrequent occasions when the files to be transferred are huge, the traditional approach of physical transfer by magnetic tapes has proved to be reasonably acceptable in most cases.

In the past few years, however, new applications profiles began to emerge. The most noticeable area is in the development of network based shared resource systems. In that mode of operation, a number of hosts or workstations share a set of common resources - primarily disc space. In this profile, the remote access of the resource is typically system initiated and as such, is done behind the back of the user. Any delay will show up to the user as irritating fluctuations in performance. Depending on the type of coupling between the workstation and the resource provider, the communication requirement can range from very severe to moderate. We have applications falling into both categories : disc-less UNIX 4.2 SUN work stations paging across the Ethernet to disc servers and IBM PC's accessing remote file servers provided by 3COM and Novell.

In the distributed processing environment CMU is creating, the main resource offered by the network is a logically "central" giant file server [1]. Work stations connected to the network will have a minimal amount of disc storage - adequate for the support of efficient system operation such as paging and file caching but not generally sufficient to hold all the files a user needs. Such file storage will be dealt with by the network file servers. This model offers a number of advantages besides cost reduction. A user can access his or her files from any station connected to the network instead of from just a specific machine; the files will be backed up regularly by a centralised professional organisation; program version control and distribution will be much easier. In order for such a distributed file system to operate effectively, a high speed local area network will be required.

Another interesting reason why the traditional serial line configuration will not gracefully satisfy the requirement of distributed processing is the fact that they have a mostly point-to-point circuit switched type of configuration. With this type of configuration, a server

machine typically has a limited number of physical access ports. Users have to contend for one of these ports in order to gain access. Due to the typically high overhead in circuit establishment, connections tend to be held for a long duration with poor port utilisation. The contention problem and circuit establishment overhead is further aggravated if a distributed transaction involves multiple servers. A possible scenario is that a station would have to go to an authentication server to get clearance before it could approach other servers for resource accesses. The probability of failure due to blocking during the transaction is, therefore, increased as a function of the number of servers it has to approach. With most local area networks, this low level type of port contention problem does not arise since they tend to operate on a datagram rather than a circuit switch model.

LAN technologies

CMU has been using a number of LAN technologies during recent years. Following is a brief summary of our experience.

Ethernet [2, 3] is a relatively mature technology and has established itself in the scientific and engineering community. Among the IEEE802 MAC (Media Access Control) standards, it has the strongest followings today and has been adopted by companies such as Digital, Xerox, SUN, TI, HP, 3COM, Interlan and Ungermann-Bass. VLSI components are readily available at reasonable cost from Intel, AMD, Fujitsu and SEEQ. CMU has been using Ethernet for over five years, starting with the experimental 3 megabit Ethernet. Three years ago, various departments, at their own initiative and funding, started installing 10 megabit Ethernets. Most of this work was not co-ordinated centrally. Currently, there are 15 separate Ethernets with over 20 segments on the campus. More than 600 high performance work stations and hosts of various types are attached to those networks.

In general, we have very favourable experience with Ethernet technology and have developed a significant amount of operational expertise across the campus [4]. One thing we have learned along the way is the need for good cable plant design. The earlier networks tended to be installed and expanded on an ad-hoc basis with little long term planning. Maintainability was usually overlooked in favour of expedient installation. Stations were attached to the most convenient or nearest location of the Ethernet trunk - often without any logging of the event. When problems arise, trouble shooting can be a long and tedious process. The simple problem of complete net outage due to shorting caused by improper transceiver installation can be determined relatively easily with the help of a time domain reflectometer (TDR). However, unless the trunk cable is carefully labelled, trying to determine the location where a certain distance, say 203 meters, from the TDR point for a given cable that meanders all over a building can be very difficult. A more common and trickier problem is the substantial performance degradation of a network caused by faulty Ethernet controllers. Typically, the controller will fail to "listen before talk" or transmit garbage packets frequently - but not continuously enough so that the jabber controller of a transceiver will shut off the station's access to the trunk. Problem determination typically involves detaching stations in a trial and error manner until the culprit is found. This is sometimes difficult when the stations are locked up in some offices. The alternative is to detach the station from the transceiver. The difficulty, in this case, is a function of the number of transceivers involved and where they are located. For a network

with over 50 stations, this can be very time consuming. We have come to the conclusion that while bus topology has the advantage of ease of installation and reduced drop cable length, star topology is substantially better from an operational and maintenance point of view. Problem determination, in that case, can be carried out at the small number of star hub locations. Given the typical lifetime of a network, the star configuration is well worth the small up front investment of longer cable runs. As a matter of fact, we have recently reconfigured a very high traffic Ethernet with over 100 stations from bus to star topology using DELNI multiport transceivers from DEC. This has reduced our typical trouble isolation and recovery time from hours to minutes.

IBM has recently introduced their star shaped token ring local area network [5]. While there are currently only interface boards available for the "standard" PC's and the RT PC's, a number of local area network vendors have expressed commitment to this IEEE 802.5 compatible offering [6]. VLSI components are available from Texas Instruments. While this network has a data rate of only 4 megabits compare with Ethernet's 10, this is not a big disadvantage. From the measurements we have done on the Ethernets, we have rarely seen network utilisation in excess of 15% - even for networks with a lot of stations relying on network file servers as well as disc servers for paging. Another potential downside of the token ring is the fact that it requires active elements in every station. This increases the potential number of points of failure. On the other hand, the token ring design has very strong operational and maintenance features built into its hardware components as well as MAC frame format and procedure. On balance, it is a good networking technology particularly for large scale LAN deployment. Furthermore, because it is topologically easier to configure than Ethernet, which is subject to the slot time propagation delay and, hence, distance constraints, the token ring is more suitable for use as the backbone inter-building network. Two rings with roughly 50 RT PC's and PC's have been operational since the beginning of 86. It has been very reliable so far. We will be increasing the deployment of this technology in the coming years.

Due to the popularity of Apple's Macintosh, particularly among students, there is a strong demand for the low cost AppleTalk [7] support. Currently, this demand has not been met. We intend to provide such services to various parts of the campus - particularly to student residences.

Over the past two years, we have been experimenting with ProNet from Proteon in order to gain insight into token ring operation. However, in view of the fact that ProNet does not conform to any standard - leaving it to be a single vendor technology, compounded

with Proteon's announced support for the IBM token ring, we do not believe we will be aggressively expanding in this direction in the future. The exception may be their 80 megabit product for backbone operation.

In general, we tend to take a dim view on any LAN approach that is not a standard and, worse proprietary. An example of this is AT&T's ISN. The implication is that network interfaces will be available from only one supplier and we don't believe that any company has the capability to produce interfaces for all the different machines in the market. This may not be such a problem in a corporate environment where equipment purchases are likely to be centrally controlled and uniform. In a university, however, every department is free to purchase its own equipment with its own grant. Therefore, we must deal with a very heterogeneous equipment environment.

In summary, from an operational and maintenance point of view, we would like to see only one LAN technology on campus; in practice, we will be required to support multiple types of LAN's. Currently, this means Ethernet, IBM token ring and AppleTalk.

Communication Protocols

Assuming we have a comprehensive physical network in place, meaningful communication between machines is still not assured since a variety of machines exists on campus and they use different protocols.

At CMU, the most popular protocol family in use is IP. It is the protocol supported by DARPA and is available under UNIX 4.2. There are over 400 stations on campus that use IP as the native protocol. These stations are attached to a variety of Ethernets, ProNets and IBM token rings. The campus IP internet is also connected to the ARPANET through an IMP operated by the Computer Science department. The second most popular network on campus is DECNET. 100 or more VAXes, Pro-350 and DEC-20 machines are connected together with Ethernet, ProNet and high speed point-to-point links. The CMU DECNET is, in turn, connected to DECNETs at Columbia, Case Western, University of Pittsburgh and Westinghouse Research Laboratory. Other protocols in use, to a lesser degree, are PUP and XNS.

From a user's point of view, the most important inter-machine communication applications are pseudo terminal access to host, file transfer and mail. We intend to provide these services between heterogenous machines. Special operating system specific communication functions such as remote procedure calls will only be considered on a case-by-case basis.

There are three approaches to solving the protocol problem. They are : (a) provide every machine with the capability to handle all other protocols in use beside its native set; (b) provide protocol translation machines; and (c) select a standard protocol and ensure all machines can handle this protocol in addition to the native set.

The first approach is quickly discounted as impractical. Protocol translation, as in approach (b), can be achieved either by implementing a set of any-to-any protocol translators or switching through an intermediate protocol. It is simple to see that the latter is preferable. Approach (c) has the advantage that, if we can find a standard protocol which has implementation on all the machines, no additional work is required.. We focused our attention on approach (c), using (b) as the fall back.

The main selection criterion for the standard protocol is that it must have strong support -

i.e. implementation available for most if not all machines. The protocol families we have short listed are : IBM's SNA, CCITT/ISO and DARPA's IP/TCP. Because of IBM, every manufacturer has tried to provide SNA interfacing capabilities for their machines. With the possible exception of mail, this protocol set is reasonably complete all the way up to the application level protocol. However, since none of the machines on the campus currently support this rather complex protocol set, we have decided against its introduction. The international standard CCITT/ISO protocol has the support, among others, of the PTT's which control all communication services in some countries. However, this protocol set is currently incomplete - particularly at the application level. While a mail protocol standard has emerged recently in the form of X400, the file transfer, access and management protocol is still pending. We have also decided against this protocol since we have no current user on the campus. The DARPA's IP/TCP protocol is very popular in the research and academic community. It is the native protocol set for Berkeley UNIX 4.2. It is a very complete set of protocols and is required for all machines destined for the U.S. Department of Defense. Due to the widespread usage of UNIX 4.2 and, hence, this protocol on the campus, we have decided to adopt it as our standard. As mentioned previously, the other popular protocol is DECNET. For our DEC-20 systems, implementation already exists that will support both DECNET and IP/TCP concurrently under TOPS-20. For the VAX-VMS systems, implementation alternatives exist. We are in the process of selecting the most cost effective approach. Implementation for this protocol is also available for the IBM PC's and Apple Macintoshes.

Interconnection of LAN's

Given that we have a number of different local area networks on campus, it will be highly desirable to interconnect them together. Physically, LAN's in different buildings can be interconnected using the large fibre optic plant that we have in place. Hence for Ethernet, we can use fibre repeaters supplied by DEC or Ungermann-Bass. However, physically connecting networks together is not very desirable since we will end up summing the traffic of the networks and, in some situations, there is the security aspect to be contended with. The ideal approach is to connect the networks together logically. We have achieved that by using the LANbridge from DEC as well as with locally developed routers.

The LANbridge is an Ethernet MAC layer (ISO level 1.5) selective relaying entity. It sits between 2 Ethernets, examines every packet on the networks and decides whether to relay it to the other side or not based on a route table. The route table is generated dynamically through observing the source address of all the packet traffic. Because of the fact the bridge has to handle and examine every packet on both networks, very high speed processing, probably with hardware assist is required - particularly in the area of table look up. The big advantage of this device is that it is higher level protocol independent. Hence it can be used to interconnect networks supporting DECNET, IP, XNS and other protocols. There are a number of small disadvantages. Currently at least, the bridge can only be used for the interconnection of Ethernets. Furthermore, each bridge can only be used to interconnect two networks. This makes the cost per connection quite high.

Another device we used for network interconnection is a router developed locally by the Computer science department [8]. It is an IP layer (ISO level 3) relaying entity.

The following is a quick tour through the algorithm :

In the DARPA world, each machine has an assigned IP address. It is a network layer (ISO level 3) address. While the physical address of the machine may change depending on the interface board used, the IP address typically remains associated with the station. In order for an IP machine (IP1) to send a packet to another IP machine (IP2), it must discover the physical address (HW2) of the recipient. If the sender does not already know the mapping, it will broadcast an Address Resolution Request (ARP request) [9]. The ARP request essentially says "I am IP1 at HW1; Will IP2 please let me know your hardware address ?" If IP2 is in the same net, it will hear the request and will reply with

its physical address in the form "Hello IP1 at HW1, I am IP2 at HW2". This method of discovering the logical to physical address mapping was a DARPA standard and is designed primary for operation within a single LAN. We extended this for multi-LAN environment. In that case, when the router hears an ARP request broadcast, it will log the fact that IP1 has a hardware address of HW1 and then will relay the request to all connected nets as "IP1 at HWR, looking for IP2". Note that it is lying to the world that IP1 has the hardware address of the router, HWR. If IP2 resides on one of the connected nets, it will reply to the router, thinking that it is IP1. The router picks up the reply "Hello IP1 at HWR, I am IP2 at HW2", logs the mapping of IP2 to HW2 and relay the reply back to IP1. The reply again will be altered as "Hello IP1 at HW1, I am IP2 at HWR". This leads IP1 to believe that IP2 resides in the router. From then on, all message from IP1 to IP2 will be addressed to the router. The same will apply to traffic from IP2 to IP1. Note that the router will only need to examine and possibly relay packets addressed to it, either directly or through broadcast. Unlike the LANbridge described earlier, it does not need to examine every packet in the networks.

The elegance of this approach is that it is simple. The original implementation was done for the PDP-11. Most of the routers currently deployed use lower cost, higher performance 68000 multibus or PC-AT based hardware. Since each router can typically support the interconnection of 3 to 4 networks, the per net cost is significantly lower than that of the LANbridge. Furthermore, the router can support the interconnection of a variety of LAN types including Ethernet, ProNet, IBM token ring, 56K and 9.6K synchronous lines. We have been using the routers for almost two years and they have been very reliable. There are three shortcomings we would like to address. First, it depends on ARP and ARP is native only to the DARPA protocol set. Hence it will not support the interconnection of DECNET or other non-IP stations. Since all machines on campus will support DARPA protocol as described earlier, it is not a serious problem. The second shortcoming is that ARP request is currently relayed as a broadcast. Since broadcast has to be handled by every attached station in the net, this can become quite expensive for a large inter-connected set of nets with thousands of stations. We have a modified version of the router under test which will heuristically relay some of the ARP requests to specific hosts instead of general broadcast. The third problem is that our current algorithm does not allow loops in the topology. This means no alternative paths for either redundancy or load sharing. While it has not been a problem for us since the reliability of the 12 deployed routers has been very high, having multiple paths supported is still desirable. We are in the process of developing a second generation router that is sub-net based. They are similar in concept to those used by MIT and Stanford and is described briefly in the next paragraph. This

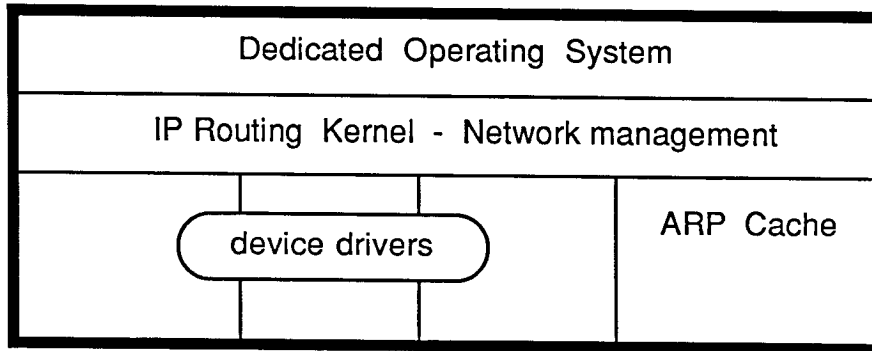
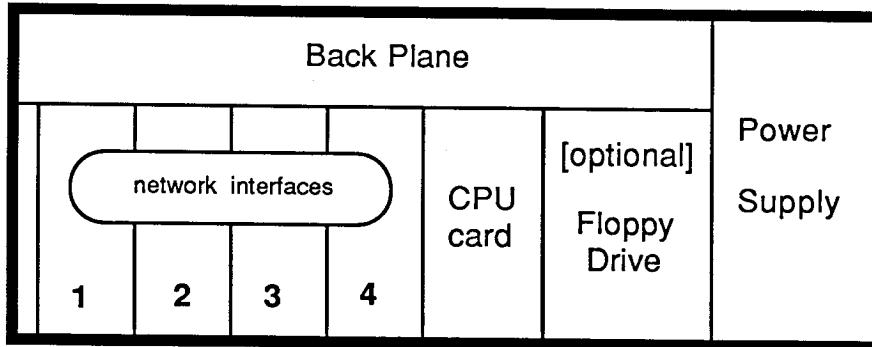
approach will address both the ARP problem as well as providing redundant paths with 'best path' selection capability.

The standard IP address is divided into a net ID part and a host ID part. If a station wishes to communicate with another station on the same IP net, it will use the ARP protocol to discover the IP to physical network address mapping. However, if the destination address resides on another net, the sending station will forward the packet to a known "gateway" instead. Gateway is described as the machine that interconnects two or more networks with different IP net ID's. The gateways will co-operate with each other to determine the best path to get from one IP net to another. In the subnet scheme, the host ID is sub-divided into a sub-net ID and a host ID - essentially extending the length of the IP net ID field. However, until recently, no standard exists for sub-net addressing and there are little or no off-the-shelf IP implementations supporting this feature. Since we do not want to change all the host software - particularly for those we do not have access to source code, we have decided against the scheme. However, our interest in this approach was revised recently when a standard for sub-netting emerged [10] and that implementation will be available with the UNIX 4.3 distribution. The potential significant reduction of ARP traffic together with the multi path capability makes the development effort worthwhile. Our design will take into consideration that during the transitional period, there will be a substantial number of stations that will not handle subnet. In those cases, the router (or gateway) will response to the ARP request based on its knowledge on how to get to the destination subnet.

Comparing LANbridge to IP router, the big advantage of the former is its high level protocol independency. By the same token, it can not address higher protocol issues such as excessive ARP broadcast nor can it support multiple paths. It is also more expensive than the router. In general, our approach is to deploy router by default. However, for networks which have a substantial amount of non-IP internet traffic also, we will consider using LANbridge. The current configuration is as shown in figure 1. The networks interconnected by bridges are those that have a large number of DECNET hosts attached.

It is our goal to provide the minimum number of "hops" between any two nets in our interconnection topology. The 10 megabit UCC Ethernet essentially fullfills the role of the backbone switch net. Most departmental networks are connected directly to this net through only one relaying element. Hence the number of hops between most networks is 2. The main exception is the Computer Science department. They operates a complex set of internal internets. In this case, only their 10 megabit public Ethernet is directly

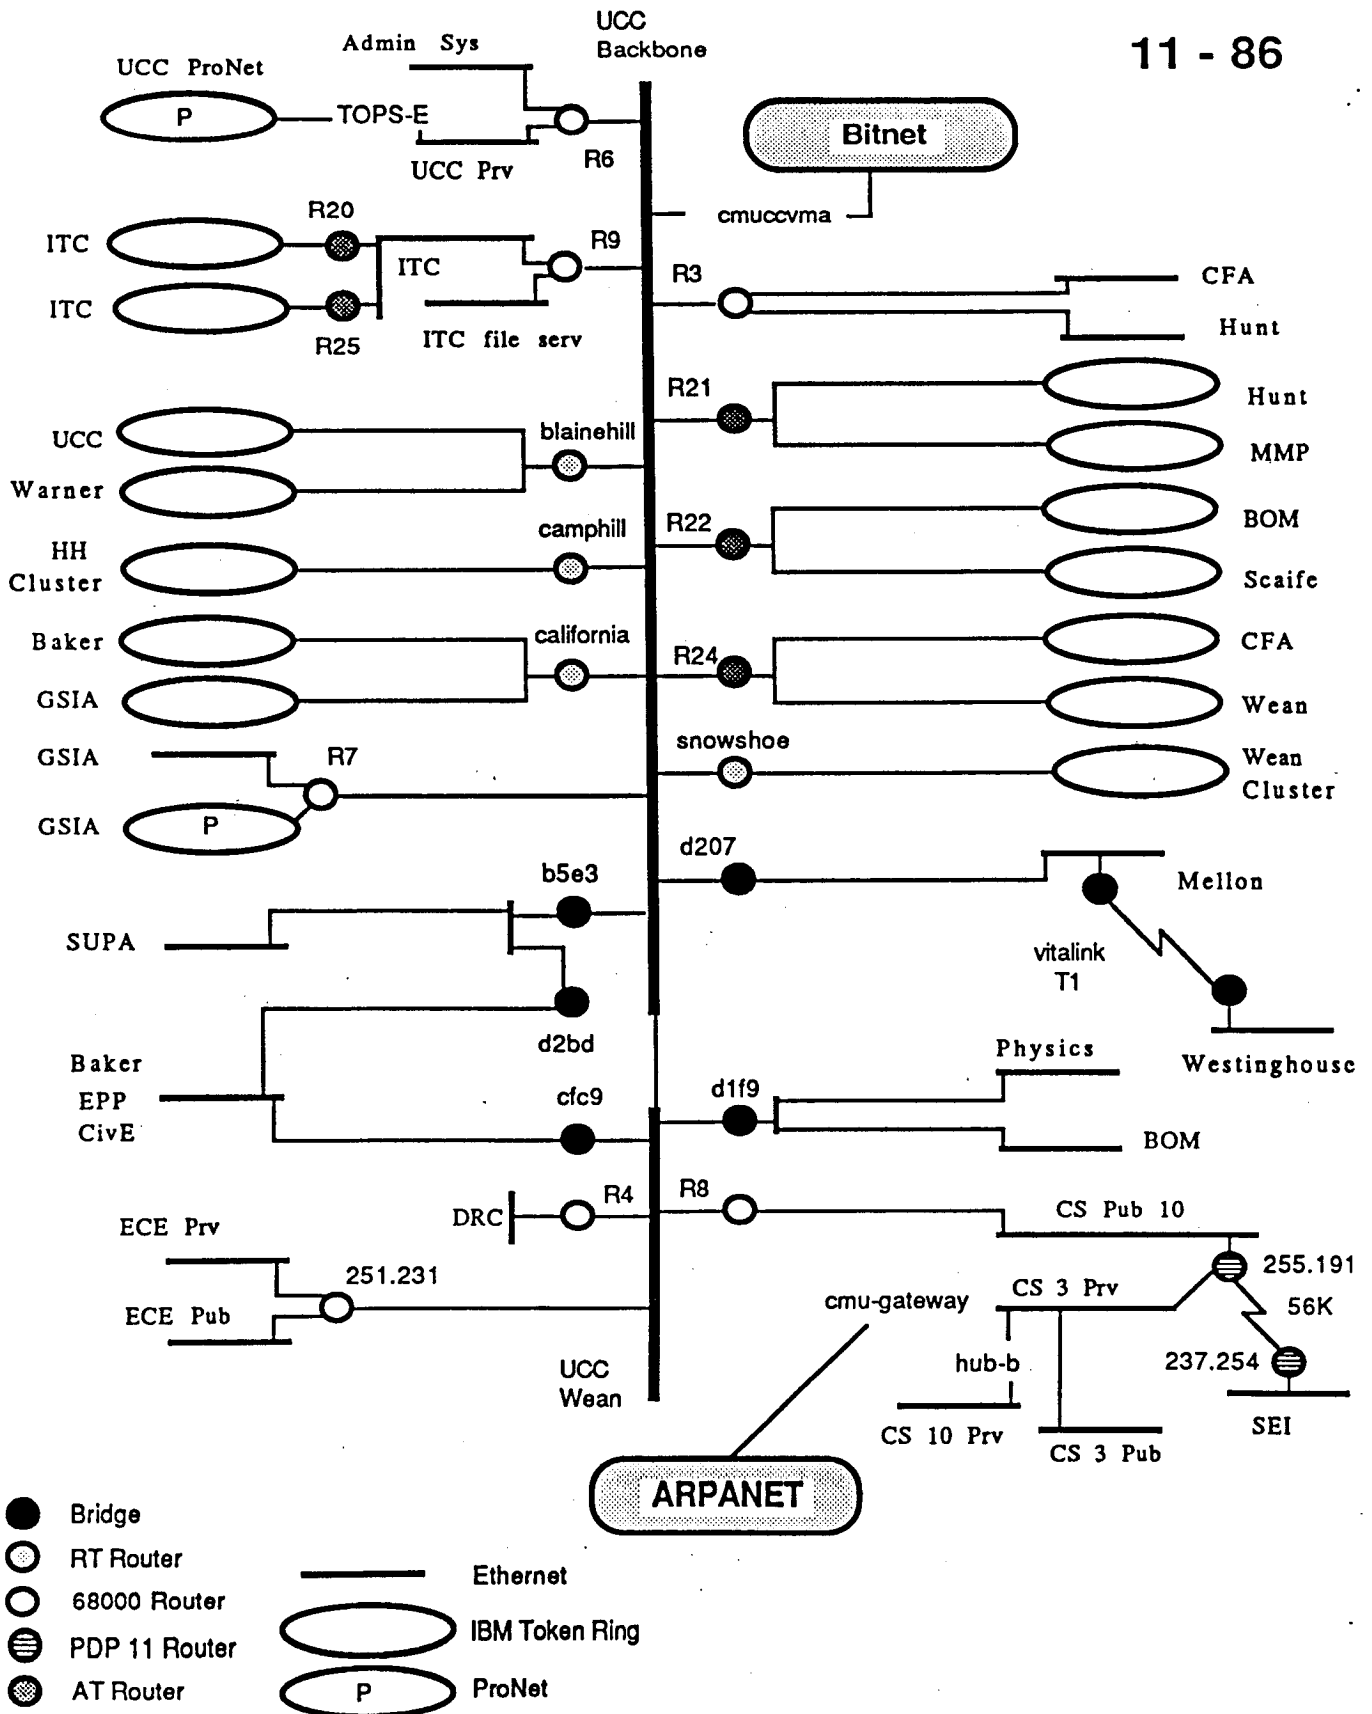
Router



PDP - 11	Unibus	CS	Ether 10, 3; ProNet; 56K syn; 9.6K sync, asyn
68000	Multibus	ECE	Ether 10; ProNet; [64K sync, IBM TR]
RT-PC	AT bus	ITC	Ether10; IBM TR
PC - AT	AT bus	ECE	Ether10; IBM TR
Kinetic	-	Com	Ethernet; AppleTalk

CMU Internet

11 - 86



connected to our backbone net. In our topology, the relaying elements are attached to the backbone rather than the departmental networks. This has the advantage of allowing the routers to be shared and hence lowering the overall cost. However, it also means that the additional inter-building cable lengths are added to the departmental Ethernets. This adds to the planning complexity for the managers of those networks. Currently, the UCC Ethernet serves both the function of a switch net as well as the main resource sharing network. Hence all the file servers are also attached to this net. When our topology is changed to one that support multiple paths, a separate server net will be installed.

The Cable Plant

While the current Ethernet, ProNet and IBM token ring cable plants have provided a large campus wide integrated high speed communication path way, expansion is still very much on an ad-hoc basis and service is by no mean universal. It is clear that this type of ad-hoc installation of Ethernet drop cables or token ring lobes will, sooner or later, lead to mass confusion. It was decided that we should comprehensively re-wire the campus - both for data communications as well as for future phone services. While this may not be the most exciting aspect of data communication technology, cable plant planning and installation, especially in case of retrofitting existing buildings, can be extremely costly. Furthermore, its impact typically spans a few generations of networking technologies. It is, therefore, most important that the cable plant selection, design and implementation be handled with proper care.

An early decision was to decouple the cable plant selection from a particular local area networking technology. It is understood that we will have to support multiple types of LAN. Today, this will include Ethernet, IBM Token Ring and AppleTalk. Over the life time of a cable plant, roughly equivalent to the life time of a building, it will be called upon to support different generations of LAN's. Hence it is very important for the cable plant to be highly adaptable.

We divided the cable plant design into two parts : the inter-buildings outside plant and the intra-building distribution system.

For the inter-buildings outside plant, the main contenders are broadband or fibre optic cables. The often quoted advantages of broadband are (a) well known technology and (b) high bandwidth, multi-channels, multi-media capability. While it is true that the CATV industry has a lot of broadband experience, the operational environment is a one way system. There is not a lot of experience on the operation of large two way communication cable plant with thousands of transmitters scattered all over the network - each capable of disrupting the complete system. While we have had some experience in the past with a small Sytek broadband network, our level of local expertise does not inspire too much confidence. As to broadband's multi-media feature, we have the following observations. For data, the high cost of broadband modems makes the media unattractive and with the exception of the IEEE802.4 token bus, most of the broadband LAN approaches are questionable. We think it is best to leave voice service to the PABX people who happen

to favor fibre optic as the main form of trunking. Broadband excels in multi-channel video which is, however, currently not a high priority requirement. On the other hand, we already have a large fibre optic plant on campus. Over the past three years, we have developed a substantial amount of expertise and positive experience with this media for data application. For outside plant operation, fibre optic's immunity to interference from lightning is a big plus. It fits well with our desire to support PABX service at a later date. It will handle point-to-point base band video if required. We have, therefore, decided to enhance our existing fibre plant instead of installing broadband as the main interbuilding trunking media. One change we will be making, however, is the type of fibre to be installed. While the current fibre plant is based on 50 micron fibre most of the data equipment manufacturers have engineered their product for 100 micron operation. The high insertion and connection loss has been a problem on a number of occasions. However, since most fibre and equipment manufacturers are converging on the new AT&T standard of 62.5 micron, we will be installing that in the future. As insurance for future expansion, we will ensure all cable runs have plenty of excess capacity for future growth - including the installation of a broadband trunk if and when the need arises.

The most complex and expensive part of the campus wiring exercise is the intra building cable plant design and installation. This is particularly true since supporting structures have to be built for all the existing and occupied buildings. Besides the high financial cost, one has to add the inconvenience experienced by the building occupants during the construction process. It is therefore, most important that we get this intra-building cable plant right once and for all.

The first decision we have to make is the cabling topology. The main candidates are bus or interconnected stars. As mentioned previously, we have come to the conclusion that star topology is preferable from a maintenance point of view. This also fits well with our telephony requirements.

The next decision is the cable type to be installed. Our alternatives are : (a) install 4 unshielded twisted pairs for phone and an additional 4 unshielded twisted pairs of data; (b) install 4 unshielded twisted pairs for phone and an additional 4 shielded twisted pair for data; (c) use the IBM composite type 2 medium which has 4 unshielded pairs for voice and 2 high quality shielded data pairs. We have decided against using unshielded phone pairs for data since it has poorer performance characteristics particularly for high speed applications. Hence it is not clear how we can support Ethernet if we install a unshielded twisted pair cable plant without also installing a parallel data network. The decision is then

reduced to whether to use the IBM composite cable or run separate shielded data pairs, specifically Ethernet drop cable, beside the phone pairs. The separate data and phone pairs approach is more costly since the cable is more expensive. We have some initial reservation with the IBM alternative. With 2 shielded twisted pairs, as available in the IBM composite cable, one can support virtually all the current LAN technologies from AT&T's ISN to token ring. With balun or transformers, it can even support coaxial based applications such as 3270 terminals and baseband full motion analogue video. In a pinch, we can even turn part of the cable plant into a broadband network. However, it cannot support Ethernet which requires 4 shielded twisted pairs for drop cable. This objection was overcome when we discovered Astra Communications produces a relatively cost effective product specifically for Ethernet operation over the IBM cabling system. Essentially, it provides an adaptor that multiplexes the collision signal onto the receive pair. We then decided to adopt the IBM cable system as our cable plant. We have also selected the actual IBM media for installation as we can obtain the material in form of a grant. Without that financial incentive, we would still have settled on the IBM cabling system but would probably have obtained the cabling media from lower cost second source suppliers.

In each building, there will be at least one wiring closet. Type 2 cable will be run from the wall outlet into the closet. All cable paths will have a fill factor of 40 percent, providing space for future expansion. The voice pair will be split out and connected to standard telephony punch block while the data pairs will be terminated in a rack mounted patch panel [11]. A typical wiring closet will support a maximum of 150 to 200 outlets. If more outlets are to be supported by a single wiring closet, the cable bundling in the area near the closet entrance can get quite severe and will require excessive amounts of cable tray or other structures. For buildings with multiple wiring closets, one of them will be designated as the Main Wiring Closet (MWC) and one as Building Entrance Facility (BEF). For most buildings, these two are actually in the same room. The BEF serves as the connection point to other buildings for voice and data using fibre cables. All wiring closets in the building will have connections to the MWC, which serves as the inter wiring closet connection center. The cable used for connection between wiring closets and MWCs will be the IBM type 1 cable which contains only the data twisted pairs. There will not be a 1:1 ratio of type 2 to type 1 cables since we are assuming that some multiplexing electronics can be placed in the wiring closets. However, in the case of voice pairs, such multiplexing will be done, if at all, only at the BEF. The phone pairs are therefore trunked from the punch blocks in the wiring closet down to the BEF using the standard 25 or 50 pair cables.

A decision we have to make during the design exercise is the number of data

IBM Cabling System

Network Independent

Star Shape

Ease of Maintenance

4 voice pairs

2 shielded twisted pairs for Data

Support :

IBM Token Ring

ProNet

ISN

Serial Line

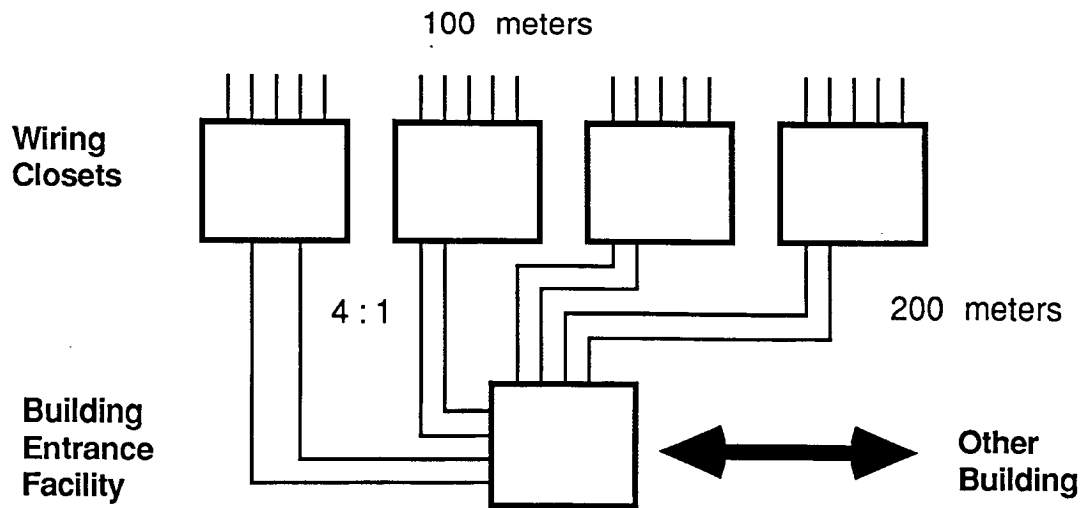
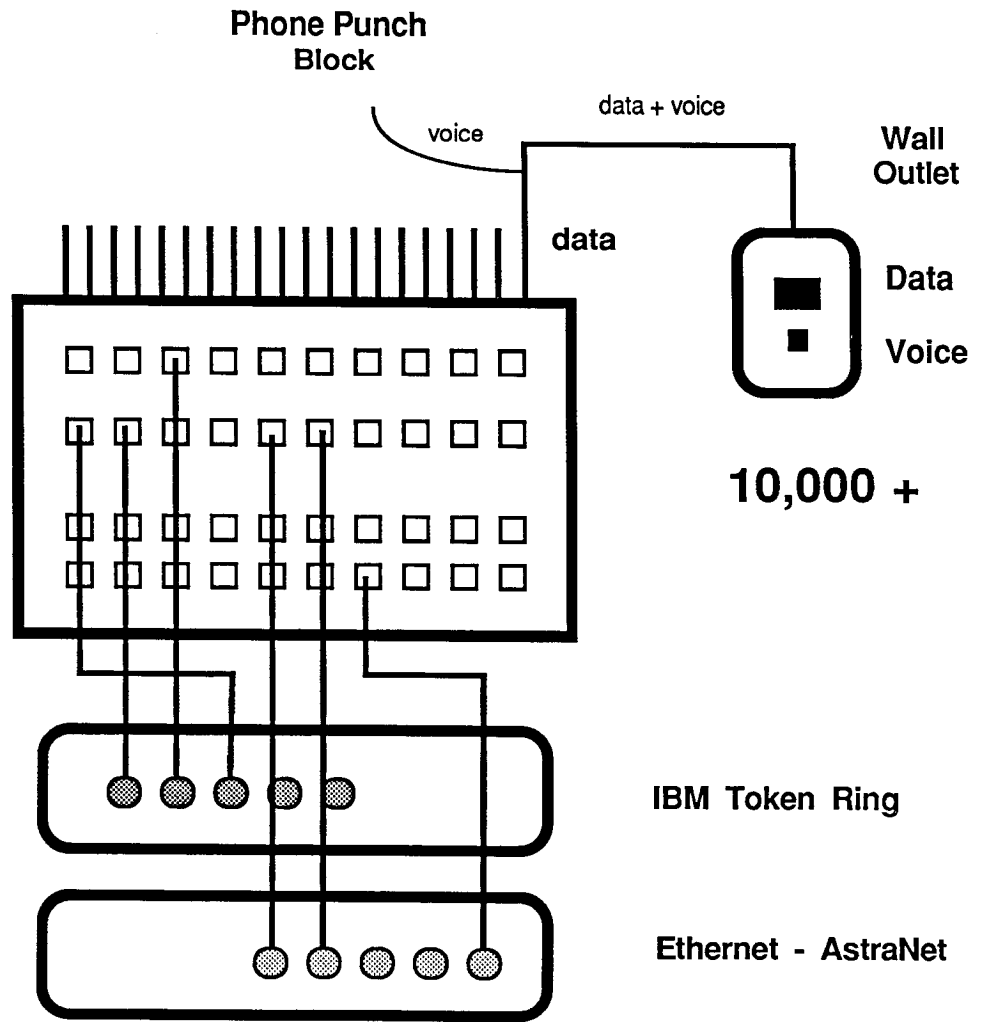
Ethernet

AppleTalk

StarLAN

3270 Coax

Full motion colour Video



communication outlets that should be installed in a room. This is a problem particularly for large rooms and work station clusters. After much deliberation, we decided to assign the number of outlets as a function of space. Additional requirements due to workstation cluster or additional occupants will be handled by local multi-port expansion modules. In the event that a large room is to be sub-divided into smaller offices at a later date, data will again be handled by local multi-port expansion units while additional voice cable will be pulled as necessary.

While we have not yet established the administration plan for the cable plant in detail, it will likely work as follows. The wiring closet will serve as a data communication bazaar supporting a limited number of "approved" networking technologies. Users will apply to the data communication office to have their outlet "activated". After registration, the cable to be activated will be connected to the appropriate networking equipment in the wiring closet. The outlet will then be clearly labelled as to what network type it has been activated for. Over and above the cable plant management, we have, under development, a comprehensive network management set of utilities. This will be the subject of a latter paper.

In summary, the size of the wiring project is very large. Approximately 50 buildings of various size, shape and vintages will be affected. Over 10,000 outlets will be installed. The current estimated cost of the project runs over \$5 million. Most of this will go into labor. The design of the cable plant is well under way. Installation has begun in the fall of 85. So far, 9 buildings have been wired. Completion is expected by end of 1986. A paper, summarising our experience will be produced upon the completion of this exercise .

References

- [1] Morris, Satyanarayanan, Connor, Howard, Rosenthal, Smith "Andrew - a Distributed Personal Computing Environment" Communication of ACM, March 1986

- [2] CSMA/CD, IEEE Std 802.3-1985 (ISO/DIS 8802/3)

- [3] Shoch, Dalal, Redell, Crane "Evolution of the Ethernet Local Computer Network" IEEE Computer, August 1982

- [4] Leong "Nut and bolt guide to Ethernet setup and interconnection" Data Communications, September 1985

- [5] Stole "A local Communications Network Based on Interconnected Token Access Rings : A Tutorial" IBM J. Res. Develop. Vol. 27. No. 5. September 1983

- [6] Token Ring Access Method and physical Layer Specifications, IEEE 802.5-1985 (ISO/DP 8802/5)

- [7] Inside AppleTalk, Apple Computer Inc.

- [8] Accetta "DARPA Internet protocol Service on the CMU Local Area networks" internal paper, Computer Science department, Carnegie-Mellon University

- [9] Plummer "An Ethernet Address Resolution Protocol" RFC826, November 1982

- [10] Mogul, Postel "Internet Standard Subnetting Procedure" RFC950, August 1985

- [11] IBM Cabling System Planning and Installation Guide, GA27-3361