The DARPA Internet Protocol Suite

Barry M. Leiner
Robert Cole
Jon Postel
David Mills

The suite of protocols supporting the DARPA internet system

THE MILITARY requirement for computer communications between heterogeneous computers on heterogeneous networks has driven the development of a standard suite of protocols to permit such communications to take place in a robust and flexible manner. These protocols support an architecture consisting of multiple packet switched networks interconnected by gateways. The DARPA experimental internet system consists of satellite, terrestrial, radio, and local networks, all interconnected through a system of gateways and a set of common protocols.

Introduction

The rapid proliferation of computers and other signal processing elements throughout the military, coupled with their need for reliable and efficient exchange of information, has driven the development of a number of computer networking technologies. The differences in both requirements and environments for these networks has resulted in different network designs. Furthermore, differences in requirements, coupled with changing technologies, has resulted in many different computer types being fielded. These different computers, although located on different networks, still have a requirement to communicate with each other.

Beginning with the ARPANET (the first packet-switched network) [1], the Defense Advanced Research Projects Agency (DARPA) has sponsored the development of a number of packet-switched networking technologies designed to provide robust and reliable computer communications. These networks have included the primarily land-line based ARPANET, packet radio networks [2,3], and satellite networks [4,5]. In addition, the use of other available technologies, such as local area networks and public data networks, has also been investigated.

As mentioned above, there is a significant need to be able to interconnect these various packet-switched networks so that computers on the various networks can communicate. Furthermore, this communication must be reliable and robust, making use of whatever communication facilities are available to accomplish end-to-end connectivity. To this end, DARPA initiated a program to investigate the issues in interconnection of different packet-switched networks. This effort has resulted in an architecture and set of protocols to accomplish this robust system of interconnected networks.

In this paper, the current status of the DARPA Experimental Internet System (the Internet for short) in terms of the architecture and set of protocols is described. The first section gives an overview of the Internet architecture, describing the key elements of the system, and their relation to each other. Following that, the set of protocols is described. Next, experiences in the test and development of the Internet are discussed. Finally, a summary and conclusions are given.

Throughout the reading of this paper, one should keep in mind that the Internet is still under development. Although a number of protocols have been standardized within the research community and are either currently Defense Department standards [4,5] or in the process of becoming standards, the Internet is constantly evolving with new functions and new protocols being developed to meet the ever-changing military requirements.

Architectural Overview

The DARPA Internet protocol suite is designed to support communication between heterogeneous hosts on heterogeneous

The views represented in this paper are those of the authors and do not necessarily represent those of DARPA, DoD, or the U.S. Government. This research has been supported under a number of DARPA contracts.

0163-6804/85/0300-0029 $01.00

networks as shown in Fig. 1. A number of packet-switched networks are interconnected with gateways. Each of these networks are assumed to be designed separately in accordance with some specific requirements and environmental considerations (for example, radio line-of-sight and local cable networks). However, it is assumed that each network is capable of accepting a packet of information (data with appropriate network headers) and delivering it to a specified destination on that particular network. It is specifically not assumed that the network guarantees delivery of the packet. Specific networks may or may not have end-to-end reliability built into them.

Thus, two hosts connected to the same network are capable of sending packets of information between them. Should two hosts on different networks wish to communicate, the source host would send packets to the appropriate gateway, which then would route each packet through the system of gateways and networks until it reaches a gateway connected to the same network as the final host. At this point, the gateway sends the packet to the destination host.

The Internet can therefore be viewed as a set of hosts and gateways interconnected by networks. Each network can act as a link between the gateways and hosts residing on it, and a gateway looks like a typical host to any network. Packets are suitably routed between the hosts and gateways to use the correct networks to traverse the system from source to destination.

Taking this view, it is clear that the service required from each network is simply the ability to carry packets between attached hosts. Gateways attach to networks as hosts. Since mechanisms must be built into the system to provide end-to-end reliability even in the face of network failures (by, for example, routing packets through alternate networks), the only service required from the network is a datagram delivery service. This means that, given a packet with a destination address on the network, the network will attempt to deliver the packet to that destination.

The overall architecture can therefore be described as four layers. At the bottom layer, individual networks and mechanisms for connecting hosts to those networks are present. At the next layer, the Internet layer, are the mechanisms for connecting the various networks and gateways into a system

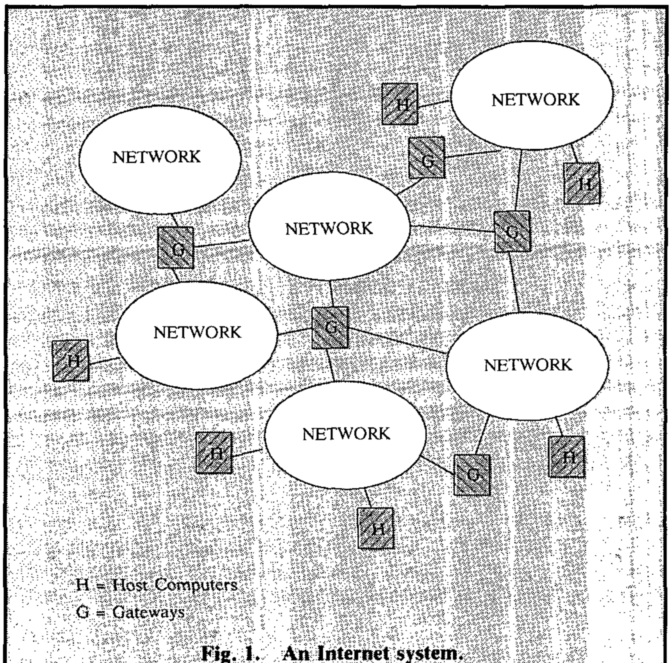

Fig. 1. An Internet system.

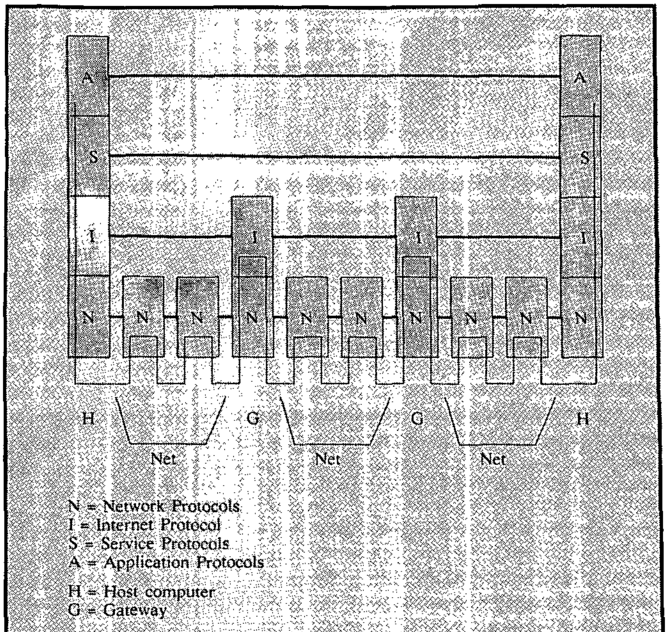

Fig. 2. The Layered Protocol Architecture.

capable of delivering packets from source to destination. At the next layer, end-to-end communication services are built in, including mechanisms such as end-to-end reliability and network control. Finally, at the upper layer, applications services are provided such as file transfer, virtual terminal, and mail.

To describe the Internet architecture, it is useful to trace a typical packet as it traverses the system from source host to destination host. Figure 2 shows the flow of a packet through the Internet. Data originates at an application layer and needs to be transported to the corresponding layer at the other end. Using the appropriate utility protocol and transport protocol, it packages the data into Internet packets. These packets are treated as data in the transmission through each of the individual networks, so that Internet packets move from host to gateway, from gateway to gateway, and from final gateway to destination host. In each case, they resemble a normal network packet on each network. The interface between the network, the hosts, and the gateways are defined by the individual networks; the hosts and gateways are responsible for packaging the Internet packets into network packets.

It should be noted that this approach, known as *encapsulation,* has some distinct advantages in the interconnection of networks. It is never necessary to build a "translation" device mapping one network protocol into another. The Internet layer provides a common language for communication between hosts and gateways, and can be treated as simple data by each network. This eliminates the "N $\times$ N" problem of building translating devices for each possible pair of networks, as it is only necessary to build the interface between the Internet layer and each individual network. Thus, hosts only need know about their local network and the Internet protocols.

The Internet Protocol Suite

To implement the above architecture, a set of protocols has been developed within the DARPA research community. These protocols have been developed with the above architecture in mind (namely a layered architecture with certain functionalities in the host-host protocols, and others in the gateways and networks). As additional functionalities have been required, either new protocols or modifications to existing ones were developed. It is anticipated that this will continue and, therefore, the description of the protocol suite given here represents the current state rather than a permanent set of "standards." Figure 3 shows the various protocols currently being used and their relation to one another.

Network Layer Protocols

At the lowest levels are the physical, link, and network protocols. These correspond to the network layer mentioned above, and provide the means for a host accessing the network. (Note that these normally describe the protocol for a host to connect to a network and not the protocol used in the network itself, that is between the switches of the network. This is of concern only to the network designer.) The key point is that the Internet accepts networks as they are, and utilizes them in an interconnected system of networks to achieve the required end-to-end communication capability. Thus, the primary areas of concern to the Internet are the interface to the network and the performance (such as throughput and delay) offered by the network.

Internet Protocol

The Internet Protocol [4,6] is the lynch pin of the Internet. It is this protocol that insulates applications programs from needing to know specifics about the networks. The Internet protocol (IP) unifies the available network services into a uniform Internet datagram service. The IP includes such functions as a global addressing structure, provision for type of service requests (to allow selection of appropriate network level services where required), and provision for fragmentation of packets and reassembly at the destination host, in the event that a packet's size is larger than the maximum packet size of the network through which the packet is about to traverse. The decision on what to put into IP and what to leave out was made on the basis of the question "Do gateways need to know it?". The key feature of IP is the Internet address, an address scheme independent of the addresses used in the particular networks used to create the Internet.

As can be seen from Fig. 2, the IP is used for communication between hosts and gateways, between gateways themselves, and between hosts on an end-to-end basis. It allows hosts to send packets through the Internet system, without regard to the network on which the destination host resides, by having the host send the packet to a gateway on the same network as the source host and letting the gateways take responsibility for determining how to deliver the packet to the final destination network (and thereby the destination host). The IP is critical to the proper operation of the Internet and the gateways in particular.

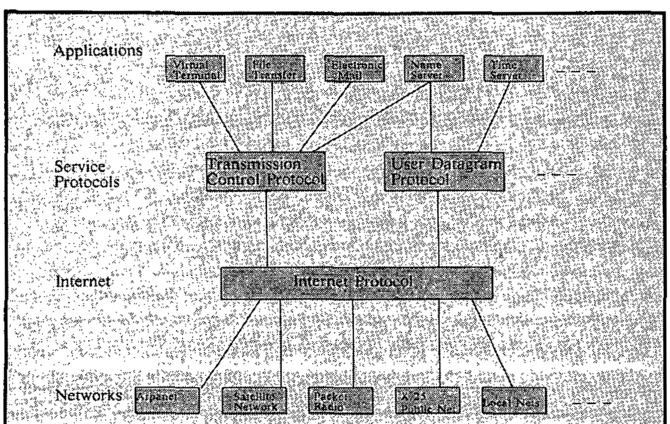

Fig. 3. The Internet Protocol Suite.

Service Protocols

The Internet protocol and layer provides an end-to-end datagram delivery service, permitting a host to inject a packet into the Internet and have it delivered with some degree of confidence to the desired destination. The customer application, however, typically requires a specific level of service. This may involve specification of reliability, error rate, and delay, or some combination of those characteristics. Rather than have each application develop its own end-to-end service protocols, it is desirable to have a number of standard services available upon which applications can build.

Currently, the DARPA experimental Internet system has two standard service protocols—the Transmission Control Protocol (TCP) [5] and the User Datagram Protocol (UDP) [8]. Other protocols are likely to be developed at this level.

In addition to end-to-end service protocols, there is a requirement for control of the Internet. An adjunct to the IP has been developed called the Internet Control Message Protocol (ICMP) [9] to serve this need.

Transmission Control Protocol

One of the prime uses for computer communication networks is the ability to reliably transmit and receive files and electronic mail. The characteristics of such use is the necessity to pass a fairly large amount of data (typically more than would fit into a single network packet) reliably and be able to reconstruct the data in sequence. To support such Internet services, the TCP was developed.

TCP provides an end-to-end reliable data stream service. It contains mechanisms to provide reliable transmission of data. These mechanisms include sequence numbers, checksums, timers, acknowledgments, and retransmission procedures. The intent of TCP is to allow the design of applications that can assume reliable, sequenced delivery of data.

User Datagram Protocol

Many applications do not require a reliable stream service. Sometimes, the basic datagram service of the Internet is sufficient for applications if enhanced by such services as multiplexing different addresses onto the same IP address and checksumming for data integrity. The UDP provides these services and permits individual datagrams to be sent between hosts. This supports applications requiring such a transaction-oriented service.

Internet Control Message Protocol

In systems as large and complex as the Internet, it is necessary to have monitoring and control capabilities, permitting hosts to interact with gateways, as well as both interacting with Internet monitoring and control centers. The ICMP provides the facility to carry out this control activity. It includes functions, such as redirect messages, to permit gateways to notify hosts that they should send packets to a different gateway (as well as error reporting).

Application Protocols

Clearly, the purpose of the Internet is to provide host-to-host and user-to-host computer communications service, thereby supporting the required applications. To accomplish this, the communicating hosts must agree on the protocol to be used for each application. A number of application protocols have been agreed upon in the DARPA experimental Internet system, ranging from the very basic terminal access protocol, to permit timesharing over the Internet through the provision of such services as name servers and time servers.

Remote Terminal Protocol

TELNET [10] is the remote terminal access protocol in the DARPA protocol suite. TELNET allows the use of a terminal on one host with a program on another host. TELNET is based on three ideas: a network virtual terminal, negotiated options, and the symmetry of processes. TELNET is built on the services provided by TCP.

The network virtual terminal idea is used to define an imaginary terminal as the standard terminal. Then all real terminals are mapped by the TELNET implementations into or out of this imaginary standard. All the data traversing the Internet in TELNET applications is in terms of the imaginary standard terminal.

The negotiated options idea calls for a base level of capability as the default operation. Enhancements may be negotiated via the exchange of requests between the two hosts. One nice feature of this mechanism is that a request can be rejected without needing to know the semantics of the request.

The symmetry of processes suggests that the TELNET protocol should work the same both ways. That the protocol is mostly used for connecting terminals to remote programs should not drive the protocol to be too specialized. It should also work to link two terminals, or to support process-to-process communication.

File Transfer Protocol

The File Transfer Protocol (FTP) [11] is based on a model of files having a few attributes, and a mechanism of commands and replies. The command and reply mechanism is used to establish the parameters for a file transfer and then to actually invoke the transfer. Like TELNET, FTP runs over TCP and thus assumes the service level provided by TCP.

Mail Transfer Protocol

An important use of computer networks is the support of electronic mail. In fact, one could attribute the success of the DARPA packet-switching research to the availability of electronic mail facilities (first over the ARPANET and then over the Internet) to the researchers involved in the effort.

The Simple Mail Transfer Protocol (SMTP) [12] is similar to the FTP protocol in that it uses the same mechanism of commands and replies. The SMTP is simpler than the FTP, in that the data exchanged is restricted to just one of the many possible combination of attributes allowed under FTP. The main concerns in the SMTP protocol are the provision for negotiating the recipients of a message, and confirming that the receiving host has taken full responsibility for the message. Like FTP, SMTP is built on TCP services.

Other Application Protocols

To illustrate some of the other application protocols that are available as part of the Internet, we describe two simple applications—a time server and a name server.

The time server [13] provides a very simple service that returns the time of day whenever it receives a request. This service may be implemented either on TCP or UDP. On TCP, if a TCP connection is opened to the server, the server sends the time of day and closes the connection. On UDP, if the server received a datagram, the server sends back a datagram carrying the time of day.

In order that users not be required to know the address of each Internet host, and to facilitate the movement of hosts to different addresses as part of normal network operations, a host name server [14] is part of the Internet. The host name-to-address hookup service is a transaction style service implemented on UDP. It expects to receive datagrams containing the name of an Internet host (for example, USC-ISIF). When such a datagram arrives, it adds the Internet address to the information and sends back a datagram carrying all that information (for example, USC-ISIF = 10.2.0.52).

Gateway Protocols

As mentioned in the architectural overview, packets flow through the system through the use of gateways located between the networks. Thus, it is necessary that the gateways communicate with each other, both for passing data packets and for accomplishing the control of the Internet, as such control is fully distributed to the gateways.

Datagrams are exchanged between networks via gateways, each of which belongs to one of several Autonomous Systems (AS). The gateways of each AS operate an Interior Gateway Protocol (IGP) in order to exchange network reachability and routing information within the AS; however, each AS may operate a different IGP suited to its architecture and operating requirements. The Gateway-Gateway Protocol (GGP), used for some time in the present Internet system, is an example of an IGP. The Exterior Gateway Protocol (EGP) is operated between selected gateways in each AS in order to exchange network reachability and routing information. Each gateway, operating EGP or an IGP, maintains a database that selects the next gateway hop on the path to each destination network. Now there are over 65 in the Internet.

All gateways support the Internet Protocol (IP) and the Internet Control Message Protocol (ICMP), which are datagram protocols requiring only minimal state storage in the gateway itself. IP support includes fragmentation, for those networks that require it, along with several options including an explicit gateway routing override for special applications. ICMP support provides notification messages to the sender in cases of misrouted traffic, excessive flows, and special maintenance messages.

Summary of Experiences

It cannot be over-emphasized that the system described in this paper is not merely a set of standards, but rather has been in operation and used on a daily basis supporting research in networking, command, and control, and other areas of computer science for over a decade. Figure 4 shows a sample indicating the breadth and heterogeneity of the Internet. The system consists of land-line networks such as the ARPANET and X.25 public networks, several phases of packet radio networks, a number of local area networks, and two different satellite packet-switched networks. Currently, roughly 100 networks and 60 gateways, all interconnected into a unified system, provide the robust and reliable computer communications service required by both military and commercial users.

The Internet has been used to support a number of applications and experiments. Interestingly enough, due to its experimental nature, perhaps its most important use in the past decade has been the support of research into networking and other computer science areas. By permitting the easy and rapid exchange of information (through both electronic mail and file transfers), as well as permitting the distributed development of software, rapid progress in these fields has been encouraged and facilitated.

The Internet has also been used to explore the implications of advanced computer communications technologies on

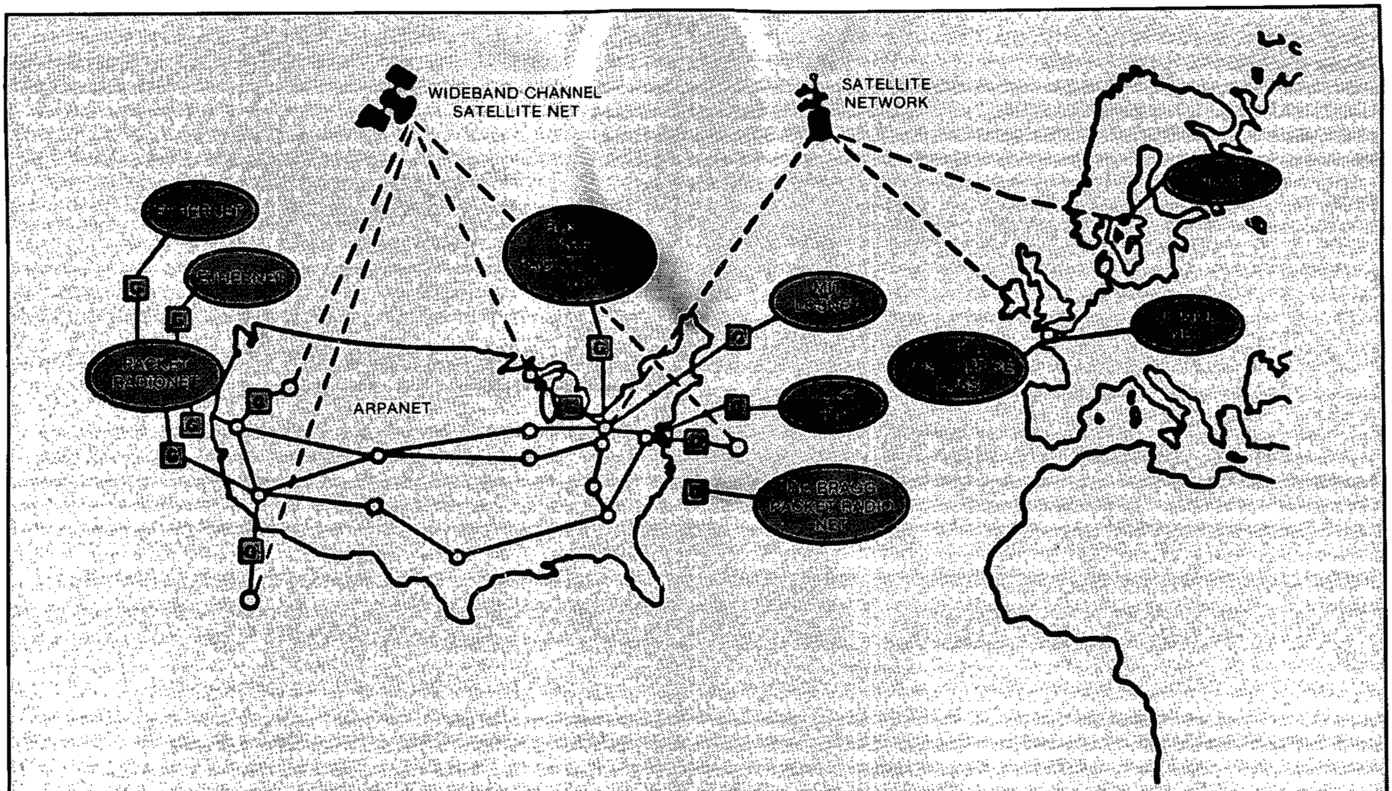

Fig. 4. The DARPA Experimental Internet System.

military concepts and doctrine. In cooperation with the U.S. Army, a testbed has been established at Ft. Bragg, NC, which is investigating the application of advanced communications and distributed processing technologies in the support of Army concepts in distributed command and control [15]. In cooperation with the Strategic Air Command (SAC), the Defense Communications Agency (DCA), and Rome Air Development Center (RADC) of the Air Force, a testbed has been established at Offutt AFB, NE, to investigate the use of the Internet technology to support strategic reconstitution efforts [15]. The Internet system provides the communications heart of a joint activity between the United States, United Kingdom, Germany, Norway, and Canada investigating command and control interoperability. In addition, a number of experiments have been carried out with the U.S. Navy using the Internet to demonstrate distributed command and control technologies. Clearly, none of these activities could have been performed with such effectiveness if it were not for the Internet system providing a unified and interoperable communication structure.

At present, the International Standards Organization (ISO) is discussing a proposal to use datagrams as the main mechanism for internetworking. The internetworking protocols will fit into a sublayer at the top of the network layer, just below the transport layer.

```
Transport layer
               / internet sublayer
Network layer <
               \ network sublayer
```

ISO has adopted X.25 as their main network sublayer protocol, and has proposed their own protocol for the transport layer [16,17]. The DARPA TCP is functionally similar to the ISO proposal for a transport protocol and can be considered equivalent.

Two groups are currently using the TCP as a transport protocol, the IP as an Internet protocol, and X.25 as the network protocol in a manner which mirrors the ISO proposals. Both groups use the X.25 network as only one component of the path between the hosts, other networks include various local area networks and the other constituent networks of the DARPA Internet.

The CSNET group uses TELENET to provide connections between a number of hosts in computer science departments throughout the U.S. [18]. The second group is a number of European research sites, the main user being University College London (UCL).

UCL provides a relay service for mail and remote login that enables U.S. and UK research workers to access each others facilities [19]. A single international X.25 connection is used to connect hosts at UCL, in England, to various Internet hosts in the U.S. [20]. The primary protocols used on the international and U.S. sides are TCP and IP. These are carried on the international public X.25 services.

Another effective use of the Internet system has turned out to be the measurement of network performance. TCP and IP can be used in network and internetwork measurements in a particularly effective manner. The protocols give two advantages:

1) The same protocols can be used over a number of networks and therefore different types of networks. This can allow comparison of network media.
2) The datagram nature of the IP layer enables network saturation measurements, while the controlled TCP allows measurement of a more conventional nature.

By using a single system to carry out measurements on very different networks, the bias due to implementation can be eliminated (for instance, in a study to compare the response to overloading in two different satellite systems [21]).

The datagram based IP enables measurements to be made of the maximum throughput a network can provide to a user.

Then using the TCP protocol, it is possible to determine how much of that throughput can be utilized by an end user, and what techniques can be used to optimize the throughput [22].

Summary and Conclusions

An experimental system and set of protocols has been described that permits communications between heterogeneous host computers on heterogeneous networks. The Internet has evolved over the past 15 years and has resulted in a set of proven and tested protocols to support the military requirements for robust and reliable computer communications. As those requirements evolve through the development of both new technologies and new military concepts and doctrine, it is anticipated that the Internet system will also continue to evolve, developing new protocols and technologies to meet those ever-changing requirements.

Acknowledgments

The Internet system has evolved over the years through the dedication and hard work of a large number of researchers. Known as the Internet Research Group, it is their efforts that made the program a success. The authors would also like to thank the members of the Internet Configuration Control Board for their many helpful comments through the preparation of this paper. Finally, we would like to acknowledge Dr. Robert E. Kahn and Dr. Vinton G. Cerf for their vision and guidance in the carrying out of the Internet research activity.

References

[1] *DARPA: A History of the Arpanet: The First Decade,* Defense Advanced Research Projects Agency, (Defense Tech. Info. Center AD A1 15440), Apr. 1981.

[2] R. E. Kahn *et al* "Advances in packet radio technology." *Proc. IEEE* pp. 1468–1496, Nov. 1978.

[3] K. Klemba *et al* Packet Radio Network Executive Summary. DARPA, 1983.

[4] Defense Communications Agency. *MIL STD 1777: Internet Protocol,* 1983.

[5] Defense Communications Agency. *MIL STD 1778: Transmission Control Protocol,* 1983.

[6] V. G. Cerf and R. E. Kahn. "A protocol for packet network intercommunication," *IEEE Transactions on Communications COM-22,* May 1974.

[7] I. M. Jacobs, *et al* "General purpose satellite networks." *Proc. IEEE,* pp. 1448–1467, Nov. 1978.

[8] J. Postel, *User Datagram Protocol,* USC Information Sciences Institute, RFC-768, 1980.

[9] J. Postel, *Internet Control Message Protocol,* USC Information Sciences Institute, RFC-792, 1981.

[10] J. Postel and J. Reynolds, *Telnet Protocol Specification,* USC Information Sciences Institute, RFC-854, 1983.

[11] J. Postel, *File Transfer Protocol,* USC Information Sciences Institute, RFC-765, 1980.

[12] J. Postel, *Simple Mail Transfer Protocol,* USC Information Sciences Institute, RFC-821, 1982.

[13] J. Postel and K. Harrenstien, *Time Protocol,* USC Information Sciences Institute, RFC-868, 1983.

[14] P. Mockapetris. Domain Names—Concepts and Facilities. USC Information Sciences Institute, RFC-882, 1983.

[15] M. Frankel, "Advanced technology testbeds for distributed, survivable command, control and communications (C3)," *Proc. MILCOM82,* paper 10.2, 1982.

[16] H. Zimmerman. "OSI reference model—the ISO model of architecture for open systems interconnection," *IEEE Transactions on Communications,* pp. 425–432, Apr. 1980.

[17] ISO/TC97/SC16/WG6. "Transport protocol specification (N1169)," *Computer Communication Review,* Oct. 1982.

[18] D. Comer. "A computer science research network CSNET: a history and status report," *Communications ACM 26.10,* pp. 747–753, Oct. 1983.

[19] R. Braden and R. Cole, "Some problems in the Interconnection of Computer Networks," *Proc ICCC* 1982, pp. 969–974, Sept. 1982.

[20] R. H. Cole, "User experience and evaluation of international X.25 services," *Proc. Business Telecom Conf.,* Mar. 1984.

[21] P. Lloyd and R. Cole, "Transport protocol performance over concatenated local area and satellite networks," *Proc. Conf. Data Networks with Satellites,* Sept. 1982.

[22] P. Lloyd and R. Cole, "A comparative study of protocol performance on the UNIVERSE and SATNET satellite systems," *Proc. Conf on Satellite and Computer Communications,* pp. 353–368, Apr. 1983.

[23] L. Palmer *et al* "SATNET packet data transmission," *COMSAT Technical Review* pp. 395–404, Spring, 1982.

Barry M. Leiner is Assistant Director of the Information Processing Techniques Office at the Defense Advanced Research Projects Agency, where he is responsible for research and development into advanced command, control, and communications technologies. Dr. Leiner received the B.E. Degree in Electrical Engineering from Rensselaer Polytechnic Institute in 1967 and the M.S. and Ph.D. Degrees in Electrical Engineering from Stanford University in 1969 and 1973, respectively.

He was a systems engineer with GTE Sylvania from 1967 to 1973. From 1973 to 1976, he was on the faculty of Georgia Tech. Between 1976 and 1980 he was a Business Area Manager at Probe Systems. Since 1980, he has been with DARPA where he has concentrated on the development of distributed communications and processing technologies.

Dr. Leiner has published papers in a number of areas including rate-distortion theory, image compression, direction finding systems, and computer communication networks. He received the Barry M. Carlton Award (IEEE AES Paper of the Year) in 1980. Dr. Leiner is a Senior Member of the IEEE and a member of Eta Kappa Nu and Tau Beta Pi.

Robert Cole is a lecturer in the Department of Computer Science, University College London, where his current research interests are in network interconnection and the communications aspects of distributed computing. He has been involved in the DARPA Internet research program since 1979 and is currently a member of the Internet Activities Board. He is the author of a text book on computer communications.

Jon Postel is a project leader at USC Information Sciences Institute. He received his B.S. and M.S. degrees in Engineering and his Ph.D. in Computer Science from the University of California, Los Angeles. He has worked at the MITRE Corporation in McLean, Virginia, and at SRI International in Menlo Park, California.

At UCLA he was involved in the development of the ARPANET Network Measurement Center and the installation of the first host on the ARPANET. Since that time, he has participated in the development of many of the higher level protocols used in the ARPANET and the ARPA-Internet. His current research interests are the interconnection of computer networks, multimachine interprocess communication, and multimedia computer mail.

David L. Mills, Ph.D., is Director—Networks, with M/A-COM Government Systems Group, Telecommunications Division (LINKABIT), and is presently leading projects in packet-switching network development and internetworking research sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA). He was formerly a Senior Research Scientist with Communications Satellite Corporation (COMSAT), working in the areas of packet-switching satellite and internetworking technologies. Before joining COMSAT, he was an Assistant Professor of Computer Science at the University of Maryland, where he was a principal investigator on several research projects in distributed computer networks and operating systems. Dr. Mills is a graduate of the University of Michigan and has held postdoctoral positions at the University of Edinburgh (Scotland) and the U.S. Defense Communications Agency. He has published and lectured extensively on data communications, computer networks, and operating systems, and has been a consultant to a number of corporations and government agencies.