Report of the ARPA/NSF Workshop on Research in Gigabit Networking

Washington DC, July 20-21, 1994

Executive Summary

In the past several years, gigabit networking research has made tremendous strides. (Gigabit networks are networks which allow individual systems to transmit data at one billion bits per second or faster). In the late 1980s, the notion of sending data at a billion bits per second over a network was an ambitious goal. It is now a reality, largely in the laboratory but increasingly in industry as well. One no longer asks if gigabit networks will become widely available but when.

A consequence of this success is that gigabit research programs need to evolve. Having demonstrated that gigabit networks are feasible, researchers must now grapple with the difficult problems of how to best use gigabit networks, how to assist in their widespread deployment, and how to ensure that networking performance continues to improve along with other computing technology. (Gigabit networking has not happened in a vacuum − rather it both reflects and has helped drive a general explosion in computer system performance).

This report was commissioned by the National Science Foundation and the Advanced Research Projects Agency and reflects the results of a two-day workshop in Washington DC on July 20-21, 1994. At that workshop, researchers from around the country were asked to present, discuss, and evaluate their views about the future of gigabit networking. This report presents what the researchers felt were the critical issues and makes recommendations about some of the issues might be addressed over the next five years.

Three general topics were surveyed at the workshop:

1. What major research problems in gigabit networking need study in the next three to five years?
2. What major research problems in applications enabled by gigabit networking need study in the next three to five years?
3. What research infrastructure is needed to make this research successful?

In addition, the workshop identified problems which, strictly speaking, are outside the area of gigabit networking but need attention if gigabit networking research is to make progress.

To help answer these issues and to focus our discussions, the participants at the workshop developed a taxonomy of gigabit problems:

1. Problems which exist only at gigabit bandwidths. These are problems such as cannot be solved without gigabit networks and applications.
2. Problems which exist at various bandwidths, but for which gigabit bandwidths require new solutions. Problems of this class are particularly common in network protocols and middleware, where the changing balance of network delays, computing performance, and bandwidth have tremendous impact on the algorithms used.
3. Problems which have the same solution at any bandwidth. Typically these problems are ones which can be solved at lower rates, but for which, if the solution is carefully designed to scale, gigabit bandwidth should not change the solution.

The workshop concentrated on the first two classes of problems, as the ones most appropriate for gigabit research, though it recognized that finding solutions to selected problems in the last category have an important impact for gigabit networking.

Networking Research. The major challenges in networking research are to take advantage of the newly developed techniques for building high-speed networks, and find ways to evolve it to meet new applications needs, to keep pace with other computing technologies (which continue to improve their performance), and to encourage the transitions of gigabit technologies into the wider technical community.

To achieve these goals the workshop identified important problems in the following areas:

1. Performance evaluation. Higher speeds and new traffic mixes are causing a much needed re-examination of models and algorithms for networking performance.
2. Switching technology. Achieving higher speeds and new types of traffic are forcing the networking community to develop innovative techniques for minimizing the cost of per packet processing in switches and routers. A continuing challenge is finding a way for these techniques to scale to switch designs with more connections per switch and higher bandwidths per connection.
3. Network management and control. A combination of new types of traffic, larger bandwidths, and the long relative delays in gigabit networks have made the problems congestion control and finding routes for data transfers substantially more difficult. Research is needed on how to best balance congestion control between the network and end-systems and on methods to quickly find valid routes for new data transfers.
4. Internetworking. Gigabit networking technologies will have to interoperate both with each other and with existing networking technologies. As a result, internetworking will be at least as important in the future as it is now. While the basic ideas of the IP architecture apply to gigabit networks it is also true that our internetworking technology needs to evolve to take advantage of the new capabilities of gigabit networks.
5. Interfacing computers and applications to networks. While it is now feasible to deliver data at gigabit rates to a computer's interface, we continue to have great difficulty getting the data through the interface and computer's operating system to the application quickly and at gigabit rates. Considerable work, probably in conjunction with the operating system community, is needed if applications are to use gigabit networks to their full potential.
6. Gigabit interfaces for PCs. Gigabit networking is no longer the domain of supercomputers and high-end workstations. PCs will soon need gigabit capabilities too and we need to encourage the development of interfaces with low costs and low heat and power consumption.
7. End-to-end protocols. Better ways to develop end-to-end protocols that meet the needs of applications are needed. Ideas like the ability for applications to synthesize new protocols from functional components need to be explored.
8. Shared media access technologies. The traditional thinking is that the high bandwidths and relatively long delays in gigabit networks limit our choices of local media access techniques but emerging research suggests that there may be a wide diversity of media access techniques that work for gigabit networks and these options should be explored.
9. Parallel channels and striping. It is often more cost effective to send data in parallel over multiple links than send the data at a higher bandwidth over a single link, a technique known as striping. While the idea of striping is well-known, it is still inadequately understood.
10. Design and verification of protocols. A tremendously frustrating problem in networking is our inability to design protocols of even modest sophistication and prove that they work correctly. Some new ideas are being developed in this area which combine design with formal verification, and given that current verification technology is nearly 15 years behind the rest of the field, we need to encourage new work in this area.

Applications Research. Now that the community has experience developing gigabit applications, a critical challenge is lowering the barriers to implementation. Currently applications are tuned to gigabit speeds on a per-application basis and there are almost no tools or techniques to assist programmers in making their applications ready for gigabit data systems. Furthermore, work needs to be done to permit even more innovative gigabit applications to be developed, and to provide security for the valuable information that gigabit applications exchange.

To address these needs, the workshop identified important research needs in the areas of:

1. Programming paradigms. Current programming paradigms for distributed systems, like remote procedure call, do not perform well over gigabit networks. It is important to either find ways to tune the performance of current paradigms or replace them with new ones suitable for higher

speed networks.

2. External data representations. It is critical for heterogeneous distributed computing to have methods for converting data between the different machine representations. Unfortunately, performing this conversion is typically a severe performance bottleneck. Better mechanisms for fast data conversion are needed.

3. Distributed control mechanisms. Gigabit networks have led to the development of new ways to perform distributed computation, such as pipelining. Unfortunately, we lack good algorithms to allow applications to manage their processing using these new techniques.

4. Adaptive applications. A new idea in distributed computing is adaptive applications, applications which change their behavior based on existing network conditions. It appears that the ability of applications to adapt will play a key role in distributed computing in the future, but we are only just beginning to understand the concept and more research with adaptive applications is needed.

Infrastructure in Support of Research. Critical to making progress in networking and applications research is the presence of a relatively large gigabit network that can be shared by applications users, applications researchers and network researchers. Gigabit networking has reached the stage where users begin to drive the research, as how they use networked applications forces applications researchers to improve their applications and network researchers to tune or even radically redesign the network to meet the users' needs. To achieve this synergy a relatively large network is needed, to create a large enough user community to help drive the applications and networking research.

This report recommends the development of a nationwide gigabit network connecting between 20 and 50 sites (where a site is typically a campus gigabit LAN) and designed eventually to serve over one thousand gigabit end-systems with several thousand users.

Needed Research in Support of Gigabit Networks. Finally, the workshop identified some research that needs to be done if gigabit networking is to achieve its full promise and recommends that efforts be made to encourage this research.

The needed research includes:

1. The development of low cost transmission technology. One of the major limitations on the deployment of gigabit networks is the high cost of optical-electrical components needed to connect to a network.

2. Faster transmission technology. While a networking topic, this area was not heavily represented at the workshop. Nonetheless, the workshop felt it was important to encourage continued research into faster transmission technologies at rates of 10 Gb/s and faster.

3. Operating systems performance. A critical goal of gigabit networking is achieving high data delivery to applications. The operating system typically stands between applications and the network (for good reason) and today's operating systems need to be revised and improved to accommodate high data rates.

4. Computer architectures. Computer architectures, in particular IO busses and memory subsystems, are often woefully inadequate for high data delivery and need improvement.

6. Security. There is concern that encryption technology is not keeping pace with progress in network performance and yet we believe that as networks get faster, more valuable and sensitive information will increasingly be carried over the network and need the protection of encryption.

7. Programming tools. Work on distributed computing over gigabit networks has forcefully reminded researchers that our programming tools for debugging distributed applications are woefully inadequate.

1. Introduction

Research on gigabit networking and the applications that use gigabit networks is reaching a new stage after nearly six years of active research. To date, the primary focus of gigabit research has been to demonstrate the feasibility of gigabit networks, show their utility for various applications and gain some initial understanding of how gigabit networks behave. This stage of research has been largely successful. Building gigabit networks is no longer considered mysterious and considerable expertise with gigabit networks has been developed in industry and academia through programs such as the six gigabit testbeds (Aurora, Blanca, CASA, MAGIC, Nectar and Vistanet). The eventual deployment of gigabit networks is seen as likely, although there remains considerable debate about when gigabits will be widely available. Pieces of industry and government are already beginning to plan on how they can best use gigabit networks when they become available.

A consequence of this success is that the gigabit research agenda needs to evolve. Simply demonstrating the feasibility of a gigabit network is no longer particularly vital. Rather research programs need to address the problems of getting gigabit networks widely deployed in a usable form and on pursuing solutions to pressing networking and applications problems that the development of gigabit networks has exposed. Identifying and addressing these problems is important for maintaining momentum towards ever improving performance in computer networking. Computer networks provide a critical service to many other fields and, to avoid becoming a bottleneck to technical growth, network performance must keep pace with the extraordinary performance improvements of computing technology in general. Furthermore, to be successful, research on performance must continue transfer successfully into real operational networks such as the Internet.

To solicit community views about how gigabits research should evolve over the next few years, the National Science Foundation and the Advanced Research Projects Agency sponsored a workshop in Washington DC on July 20-21, 1994. The invited participants in the workshop were experts in various areas of gigabit networking and applications. The workshop was charged with producing a report that answered three major questions:

1. What major research problems in gigabit networking need study in the next three to five years?
2. What major research problems in applications enabled by gigabit networking need study in the next three to five years?
3. What research infrastructure is needed to make this research successful?

ARPA and NSF encouraged the participants to consider these research questions broadly. Researchers were asked to consider research questions outside their particular sub-disciplines of expertise and, further, to note cases where gigabit networking and applications depended on research by other communities. Government attendees were urged to speak as technical experts rather than as leaders of various government programs. The meeting was structured as a mix of breakout and plenary sessions designed to encourage exchanges of ideas across sub-disciplines. This report is the product of workshop and summarizes the recommendations of the participants (who are listed in Appendix A).

1.1. The Role of Gigabit Networking

Gigabit networking activities are an important part of a larger pattern in technology. Technology trends continue to deliver impressive improvements in performance and reductions in cost. For instance, computing power for a given price roughly doubles every eighteen months. At the same time an increasingly large amount of information is being stored, manipulated and moved in digital form. Sound and video, both once reliant on analog transmission techniques, are now handled by digital systems.

These trends are an important motivation for gigabit networking. Computers are capable of handling ever larger amounts of data at ever faster rates and are therefore placing demands on networks to transmit and deliver data at ever high speeds. The digitization of sound and video is challenging the current division of transmission technologies into separate telecommunications, cable TV and data communications communities because there may no longer be a need for different types of infrastructures to handle analog signals. All telecommunications and data communications networks are or will soon be transmitting entirely digital data and cable TV networks are also moving in this direction, which permits one to consider using a single transmission infrastructure to support all three industries.

At the same time these trends pose tremendous challenges. Until recently improvements in networking performance have lagged behind trends in other computing technology and it remains a serious technical challenge to force networking technology to keep pace. The potential use of a single transmission technology for voice, video and data requires a versatile transmission system capable of meeting the quality of service (QoS) requirements of the different technologies. And the research community needs to find effective means of technology transfer, so that technical advances move smoothly from the lab, to experimentation over experimental high-speed network infrastructure, to wide use on general networks.

1.2. Distinguishing Gigabit Problems

In the initial phases of gigabit research, research problems were judged by whether they uniquely required gigabit-per-second performance. Problems which failed this gigabit litmus test were deemed unsuitable for investigation under the auspices of gigabit research. Given that the nature of gigabit research is clearly changing, some time at the workshop was devoted to trying to understand what the suitable litmus test for future gigabit research should be.

An important concern was that the strict rule that a problem must require gigabit performance excludes a certain class of research that must be done before gigabit networks can be widely deployed. While using such a strict rule as gigabit research was first getting started was essential to focusing research, as gigabit issues have become better understood, it has become clear that certain problems can be solved in lower speed networks and then scaled cleanly to gigabit networks. One example in networking protocols is the development of Asynchronous Transfer Mode (ATM), a ‘‘gigabit protocol’’ which was designed to scale cleanly from 64 Kb/s speeds to 10s of Gb/s, and is largely being tested at speeds of a few hundred Mb/s.

After some discussion, the workshop generally settled on a three part classification system for research problems:

1. Problems which exist only at gigabit bandwidths. These are problems such as real-time distributed databased access, high-quality conferencing and real-time distributed manufacturing, as well as many Grand Challenge problems, which cannot be solved without gigabit networks and applications.
2. Problems which exist at various bandwidths, but for which gigabit bandwidths require new solutions. Problems of this class are particular common in network protocols and middleware, where the changing balance of network delays, computing performance, and bandwidth have tremendous impact on the algorithms used.
3. Problems which have the same solution at any bandwidth. Typically these problems are ones which need to be solved, but for which, if the solution is carefully designed to scale, gigabit bandwidth should not change the solution. An example of this class of problem might be self-authenticating messages. If they can be authenticated at low bandwidths they can presumably be authenticated just as well at high bandwidths.

The workshop concentrated on the first two classes of problems, as the ones most appropriate for gigabit research, though it recognized that finding solutions to selected problems in the last category have an important impact for gigabit networking.

1.3. Outline of the Rest of the Report

The rest of this report presents the conclusions of the workshop.

Section 2 presents critical research problems in gigabit networking that the workshop identified.

Section 3 presents critical problems in gigabit applications.

Section 4 presents the workshop’s opinions on what kinds of research infrastructures are most suitable to stimulating and supporting gigabit research in applications and networks over the next three to five years.

Section 5 identifies related research problems. Over the course of the workshop it was clear that just as many disciplines increasingly are relying on high speed networks, so too high speed networking research depends on progress in other fields. Section 5 identifies some research areas in which progress is needed if gigabit research is continue smoothly.

2. Gigabit Networking Research

One of the three research areas given detailed attention at the workshop was gigabit networking research − research into protocols and services built over gigabit networks.

2.1. Introduction

Gigabit networking represents an attempted evolution of at least relatively mature fields (telecommunications and data communications) and possibly a third (cable television). Each of those fields has its own design principles and goals, which have been slowly coalescing together into a single set of objectives for gigabit network. This section summarizes what we believe are the joint objectives.

2.1.1. Improving Networking Performance

Gigabit networking research activities seek to develop a fundamental understanding of the key technical challenges that must be met in order to dramatically improve the performance of networking technology, while providing the functionality needed for general use and achieving costs that are low enough to stimulate wide-scale deployment. This work seeks to use this understanding to create innovative technical solutions and demonstrate the effectiveness of these solutions through the construction of experimental systems. Research can most directly benefit the nation as a whole if it is directed towards providing the technical leadership for the development of the National Information Infrastructure. A key objective for NII development is the creation of an open architecture that actively encourages continuing innovation and facilitates participation in its development and use by a broad spectrum of individuals and organizations. The leadership of the networking research community is essential to achieving a truly open architecture and standards that maximize interoperability, ensuring a broad range of choices for end users.

While the word ‘‘gigabits’’ as been used to identify a particular suite of gigabit networking research, the objective is not just to demonstrate operation at any particular speed. The fundamental objective is to accomplish a substantial increase in the effective speed of the delivered service, in the context of the performance of the underlying technology at any given time. That is, given a set of assumptions about the performance and cost of digital logic, lasers, and so on, can we identify changes and innovations in network organization, architecture and function that permit a substantial increase in delivered throughput. Clearly, one could passively accomplish gigabit operation by waiting a decade for improvements in the underlying circuit technology. However, if we can combine those anticipated improvements with architectural and algorithmic innovations in the network design, we can obtain a major step forward in performance.

Once such a step forward is achieved, the goal is then to maintain the new level of relative performance as technology continues to improve. Much as computer processor technology can be expected to offer doubled performance very eighteen months, networking technology has to be designed to scale so that it too can show continued performance improvements as technology trends dictate. As evidence of the pressure of technology trends, it is worth noting that when gigabit research began in earnest six years ago, gigabit networks were viewed as primarily the domain of supercomputers. A few years from now, PCs will both use, and quite probably need, gigabit networks. Observe that one of the challenges is to find technologies that gracefully scale over various speeds, because it is far easier to replace an endpoint in a network (a single computer) than it is to replace an entire network.

2.1.2. The Tension Between Speed and Flexibility

As we move to higher speeds for a constant performance of the technology, we are forced, of necessity, to reduce the baseline requirements on that technology. For example, there is processing required in the switching of packets or cells, to accomplish routing, quality of service, accounting, security, and so on. In the early days of packet switching, these functions were performed in software, using a general purpose processor. While the speed of processors clearly increase with time, at any point that performance bounds the achievable limits on packet processing rates. If we wish to increase the processing rates for a given generation of technology, we can either abandon the general processor for more specialized and presumably more efficient (but possibly less flexible) alternatives such as specialized hardware, or make more efficient use of general processors by reducing the required functions. In both cases, there is a pressure to simplify or at least regularize the processing. This tendency, however, may be in conflict with other needs in the network. Complex application requirements, for example multiple qualities of service, may imply sophisticated processing that is hard to reconcile with the objective of simplicity. Further, the application requirement will evolve over time, which implies a need for these simple functions to change, or at least to be combined in new ways. We can thus state, at a high level, a key research objective at the network level, which is to find ways to make both simple and flexible use of the underlying technology, with the objective of achieving both a substantial speed increase and a sufficient flexibility to meet a range of application requirements.

2.1.3. Latency Issues

Although computers and transmission links increase in speed, the speed of light, and thus the round-trip latency of network links remains constant. This fact represents a key challenge to networks, since from the perspective of the processor speed, the latency is becoming ever longer. The roughly 100 ms round trip across the United States, which a decade ago might have stalled a process for 100,000 instructions, now stalls it for 10 million instructions. Increasing processor speeds will only make this worse. Increasing transmission speeds do not eliminate this problem, but only reduce the impact of sending large messages rather than small ones. Since this limit is fundamental, a general architectural principle, which is well understood, is to reduce to the extent possible, the need for round-trip interactions in the network. For example, an application can send one complex query across the network, instead of a sequence of simple ones. More broadly, this points to a general need for developing methods of distributed parallel processing that can achieve high performance using either ‘‘coarse-grained’’ interactions or careful highly pipelining of data flows from one system to the next.

2.1.4. State, Simplicity and Latency

One way to simplify the requirements for packet processing is to establish some state in the switches which describes what actions to take when a packet arrives. This distinction, which has sometimes been called the connectionist vs. connection-less processing tradeoff, attempts to reduce the frequency with which complex decisions are made, so that the processing of individual pieces of data within a data stream, can be streamlined. Thus, the circuit setup phase of ATM (or other connection-oriented networks) establishes per-connection state, and pre-defines what processing must be done for each of the data elements passing along the virtual circuit. As we move to more complex application functions, such as explicit Quality of Service, the need for this state likely increases. However, this approach of pre-installing state must be balanced against the need to reduce the frequency of round-trip interactions.

The tension between the need for state and the need to reduce setup latency can be resolved in a number of ways. First, the state, once established, can be used for a number of interactions, so that only the first unit of data is delayed, assuming the state is relatively stable. Second we can use caching techniques, in which state information is retained on the chance it can be used again. Third, we can use techniques that attempt to establish state ‘‘just-in-time,’’ by preceding the data with a state establishment message, but without demanding a round-trip co-ordination in the establishment of the state. All three techniques can apparently be usefully applied in high speed networks, although the balance among approaches (which is determined by the processing required to create the state initially, the cost of retaining possibly outdated state, and the frequency with which state is required to change) is not yet fully understood.

There is also an emerging school of thought that suggests that state should be created and managed in such a way that the temporary loss of state (for instance, due to a brief switch or router outage) produces possibly degraded performance rather than a failure of connectivity. Use of such ‘‘soft state’’ has the potential to put more state information into the network while encouraging flexibility and robustness of the network.

2.1.5. Dynamic Range and Heterogeneity

The existence of high speed networks does not eliminate low speed networks. Heterogeneity will always be present, both in the technology options (e.g. Ethernet and Token Ring as current generation LANs) and in the speed options which must be accommodated in the network. A key issue for gigabit networks must thus be to permit the interworking of both high speed and low speed components. This ability to adapt must be present both in applications and in network level mechanisms. The current approaches to heterogeneous interworking, such as the Internet protocol suite, are based on a set of assumptions about the required range of adaptation to the technology below and the applications above. Both of the assumptions are being stressed at present. Higher speed networks, with the implications discussed above for simplification and (possibly) for increased dependence on state in the network, challenge the simple assumptions of the Internet model. At the same time, more complex transmission capabilities (and their use by applications) place more constraints on our ability to interwork among different technologies. Heterogeneous interworking, while not just a matter of speed, is a major issue in gigabit networks, since we must assume that gigabit networks, like all networks before them, will contain heterogeneous components, as well as the broad range of speed options from the advanced high speed nets to the legacy links of the past.

2.1.6. Cost and Performance

Our objective is not just to find theoretical solutions, but to demonstrate practical operation in real networks. As an engineering discipline, we cannot ignore the real issue of cost-effective solutions. In particular, it is important to push gigabit networks into the workstation and PC environments, and not just the supercomputer center. To accomplish this objective, it is important to devise gigabit approaches whose costs are consistent with the cost profiles of advanced workstations and PCs of today. We can assume that, with some prodding, network and workstations costs can be made to track as they evolve with time. Thus, if we can find architectural approaches which provide a step improvement in workstation network performance today, we can expect to maintain this performance relationship into the future. Thus, this gigabit research initiatives should include a focus on the objective of identifying and demonstrating the key technology and architecture issues that will advance network performance in concert with common computing technology like workstations and PCs.

While recent years have seen striking improvements in our understanding of the fundamental issues that determine the cost and performance of competing network technologies, our debates on these issues are often too little informed by a quantitative assessment of how cost and performance relate to one another. The situation is not unlike that of computer architects in the early eighties when the debate over RISC vs CISC architectures first emerged. In the intervening years, computer architects have developed a systematic methodology for evaluating architectural alternatives that uses a combination of analytical techniques and trace-driven simulations derived from standard benchmarks, to evaluate different approaches, while accounting for both cost and performance. Networking research too often fails to evaluate cost and performance in combination with one another, treating them as independent, rather than closely coupled. Achieving significant improvements in networking will depend increasingly on our ability to quantify the impact of architectural features on both performance and cost under realistic operating conditions.

2.1.7. Summary of Objectives

On the basis of the preceding discussion, we can identify a set of fundamental objectives for high speed networking research.

1. Make the in-band processing of data in the network very simple and regular, with the objective of a step performance increase relative to the underlying technology.

2 Perform this processing in ways that can evolve to meet new application needs and can effectively exploit continuing improvements in the underlying technology and keep pace with other computing technology.

3. To the extent that state is established and preserved inside the network to simplify this in-band processing, minimize the frequency with which state-establishment requires round-trip delays.

4. Create strategies for incorporating successive generations of high performance technology into a heterogeneous networking infrastructure and determine how to apply high performance networking principles to scale up performance at all levels, including internetworking and end-to-end transport.

5. Develop more systematic methods of evaluating the cost and performance of alternative technologies and mechanisms and apply these methods to assess competing approaches.

In the remainder of this section, we discuss a number of specific technical issues, which bear on the problem of improving the throughput performance of networks.

2.2. Gigabit Networking Research Topics

This section attempts to ennumerate the key gigabit networking research problems that the workshop believed need to be tackled in the next five years. This list, of course, is imperfect. Research results in the coming will uncover critical problems requiring attention. Similarly it may prove that certain problems can be finessed and no longer require attention. This list simply represents our best understanding of the state of the research problems today.

2.2.1. Methodologies for Performance Evaluation

High-speed networking has encouraged the development of new switching technologies as well as the development of real-time throughput-intensive distributed applications that were not feasible at lower speeds. To provide the services required for these new applications, high-speed networks must be carefully managed. Network management and control includes network design, routing, and flow control, each of which operates at a different time scale. Models of user traffic as well as network behavior in response to that traffic form the basis for each of these network control procedures.

In lower-speed networks, short timescale traffic control algorithms, such as routing and flow control, are usually designed to be reactive. However, in high-speed networks, predictive rather than reactive traffic control is likely required since feedback delay can exceed the duration of a traffic event in the network. Predictive traffic control requires accurate models of user traffic as well as network behavior in response to that traffic. However, for high-speed networks, the behavior of many new applications and the response of the network to the application traffic is not yet well-understood.

Analysis of application traffic through both detailed simulation and modeling of the application as well as measurement of actual traffic is needed to provide accurate characterization of application behavior. This problem has recently become more urgent as new research has shown that data traffic apparently has very different characteristics from voice traffic and may require us to discard or completely rework our analytical traffic models. The goal of this analysis is to produce a small set of simple traffic models that represent the variety of new applications developed for high-speed networks. These models will be used as input to the traffic control procedures as well as to the network design algorithms.

Analysis of the behavior of individual network components, as well as collective network behavior in response to both actual and modeled application traffic is needed for network design and for the development of effective traffic control algorithms for high-speed networks. As with the applications, a combination of simulation, modeling, and actual measurements should be used to characterize the behavior of the switches and the network as a whole.

Network designers also need better mechanisms for relating performance to cost. The common practice of decoupling network performance evaluation from cost considerations prevents meaningful assessment of alternatives. Improving the state of the art in performance evaluation requires new tools that can incorporate the characteristics of real traffic (possibly something akin to the trace-driven simulations, derived from computational benchmarks used by computer architects) and allow side-by-side comparisons of competing architectural approaches as well as detailed designs/implementations. Analytical and simulation tools are needed for evaluating decisions in advance of constructing systems, while measurement tools are needed for evaluating actual systems, tuning their performance and obtaining the data needed to drive simulations.

2.2.2. Switching Technology

Current packet switching technology is based on simple approaches to cross-connection, most commonly the use of shared memory to provide both buffering and cross-connect. A key issue in increasing network throughput relative to technology is to explore switching options that exploit parallelism to allow the construction of scalable larger switch architectures or scalable networks of many small switches.

Scalable architectures are essential to providing large networks that deliver high performance to every user while maintaining a low incremental cost per user. While there is a fairly good understanding of how to design such systems to provide scalable, point-to-point communication at a low cost per user, we have a much less complete understanding of how to provide low cost switching of multipoint (one-to-many, many-to-many) channels. This is particularly true, when one considers the need for reliable multipoint communication (rather than simply best-effort) that is inherent in many high performance, distributed computing applications.

The current ATM technology has simplified the per data unit processing within the network by standardizing on a single fixed size cell. While the use of a fixed size cell offers certain advantages, it is not without its drawbacks. Direct switching of variable size packets can reduce the amount of network bandwidth devoted to carrying overhead information and can simplify (or eliminate) the adaptation layer processing currently needed in ATM. The switching of larger data units allows more time for decision making on each data unit, making it easier to make the per data unit decisions at high speed and allowing more complex decisions (such as datagram routing) to be performed on each data unit. While ATM will clearly play an important role in high performance networks, other alternatives should be explored in order to develop new paradigms for successive generations of gigabit networks.

Virtual circuit networks, such as ATM, provide a separation of the per data unit processing from the control operations needed to enable the efficient processing of user data streams. While the data path of the switching systems determines the performance users perceive once communication channels have been established, the design of the switching system's control subsystem determines the speed with which it responds to user requests for channels and the frequency with which the characteristics of those channels can be modified. The design of the control subsystem is thus central to obtaining good performance in applications with highly dynamic communications requirements. The issues involved in the design of high performance, scalable control systems for large virtual circuit networks, supporting multipoint communication have received only limited attention to date. This remains an important challenge for switching systems researchers.

2.2.3. Network Management and Control

Signaling provides a means by which applications communicate control information to the network and by which different network components communicate with one another, in order to implement the applications' requirements. Signaling can be employed to request a specific quality of service from the network, or to establish state for the purpose of simplifying the in-band processing and thus improve overall performance. Signaling methods vary widely among the different existing networks, from the very simple interfaces of the connectionless data networks to the complex signaling of the phone system. As networks speeds increase, the use of signaling to communicate performance requirements to the network should allow the network to make better choices on behalf of applications, improving both their individual performance and the efficiency of network utilization. As networks and applications become more complex and heterogeneous, signaling systems can provide a stable common interface among varied switching hardware at the lower levels and varied application needs, up above, enhancing interoperability and making it easier to upgrade system components with improvements in technology. Signaling can also provide the means to identify the need for format translations or other operations, to provide seamless end-to-end communication. Finally, as distributed applications require more extensive use of multipoint communication channels, signaling mechanisms for creating and maintaining those channels become essential.

Routing can be usefully divided into two parts. The first is concerned with acquiring information about the state of the network (how busy its links are and which links are up or down), while the second is concerned with using this information to select paths between two participants requiring point-to-point communication, or the tree joining a collection of participants requiring multipoint communication. The challenge in gathering state information is to determine what information is really needed to make intelligent routing decisions (routing decisions that satisfy users' requirements, while maintaining efficient network operation) and the frequency with which it must be updated. The effort required to obtain complete information on a continuous basis may not be justified by the marginal performance improvement over schemes that use more limited information gathering. In gigabit networks, the type of information needed for routing may be different, the large scale of networks may limit the frequency and scope of the information that can be usefully acquired, while the speed of the networks themselves may reduce the cost of maintaining the information. The path (or tree) selection portion of the routing process, is typically performance-critical, since the speed of routing decisions affects either the time required to deliver a piece of data, or the time required to establish (or extend) a virtual circuit. At the same time, the nature of the routing decisions can be complex, since they may need to consider quality of service requirements and the potential impact of a routing decision on other users. Determining the best topology for a multipoint communication channel is computationally challenging, particularly under the common practical constraints of distributed routing algorithms with limited information, and dynamically changing sets of participants.

The heterogeneity of link speeds and the speed of application streams contributes to the potential for congestion in gigabit networks. Mechanisms are needed both to smooth the discontinuities introduced by different link speeds and to ensure high throughput and fair treatment of different streams during those periods when the network is overloaded. Gigabit networks raise new challenges for congestion control since they involve a much larger bandwidth-delay product than earlier network technologies and permit a wider diversity of bandwidths in paths through a network (e.g., one path may include both Gb/s and 9.6 Kb/s links), both of which reduce the effectiveness of some classical congestion control techniques. The high performance requirements limit the time available to make decisions, which together with the objective of low cost implementations creates a pressure to carry out only the simplest kind of processing on a per data unit basis. The problem of guaranteeing quality of service to applications in the face of unpredictable behavior on the part of at least some applications, remains only partially understood. Finally, we need to determine how best to divide the responsibility for congestion control between the network and the end systems.

2.2.4. Internetworking

Successful gigabit networking technologies will have to exist in a complex environment that includes networks owned and administered by a wide range of different entities and implemented using a variety of base technologies, including both current legacy networks and new technologies that will be created in the future. While there is a continuing temptation with each new generation of technology to think that it will take over and replace the prior generation, experience shows that new technologies only increase the level of heterogeneity. They never seem to decrease it. Hence, internetworking is an essential characteristic of gigabit networks and must be a central component of a gigabit networking research program.

Many basic features of the IP internet model have enduring relevance in higher speed networks, but it is also clear that the new subnetwork technologies and new applications place demands upon the internet layer that are not currently being met. As with the underlying subnets, the internet layer can profit from simplification of the per data unit decision making. By reducing the frequency with which complex decisions must be made, the internet layer can streamline the data path, process data more quickly (possibly moving some operations into dedicated hardware) and can insert operations that provide ‘‘matching’’ between incompatible subnets on an as-needed basis. Applications’ needs for specified quality of service have an impact on the internet layer, in addition to the subnets and the applications. The increasing demands being made by applications makes it more important that the capabilities of underlying subnets be fully exploited, meaning that the internet layer may need to make more sophisticated decisions when selecting the subnet used to support a given application. Multipoint applications, such as audio and video distribution, which are now being supported as an internet overlay, should probably be supported more directly at the internet layer.

An important aspect of the internet layer is that it lies at the neck of an hourglass shaped architecture. Below it are a myriad of transmission technologies. Above it are a diversity of applications. If the applications are to exploit the power of improved transmission technologies, the selection of the set of functions to be placed in the internet layer must be chosen with care. In addition, high performance scalable switching systems for internet router applications are clearly going to be needed to enable the internet to cost-effectively support large numbers of end systems communicating at gigabit rates.

2.2.5. Computer-Network Interface and Gigabit End-Working

While emerging network technologies are capable of delivering data at very high rates to end systems, applications rarely see more than a fraction of the throughput the network is able to deliver. This reduction in performance is due to a combination of factors, including limited effective memory bandwidth, lack of general purpose mechanisms for moving data directly among peripherals without the intervention of software processes in the data path, the organization of the memory hierarchy (cache, main memory, virtual memory), operating system scheduling, context switch overheads and the communications interface to application programs.

Applications requiring both continuous data streams (such as multimedia) with stringent quality of service requirements and gigabit data streams that may be highly bursty, place new requirements on both the hardware interface between the computer and the network and the software components that mediate between the hardware and individual application programs or devices. To enable higher performance to applications, there needs to be a more direct coupling between the network and the applications and better allocation of processor resources to meet QoS requirements and ensure fair treatment during overloads. At the same time, it's important not to sacrifice the protection that modern operating systems afford to applications. Indeed, new applications will place ever more stringent requirements for security and protection on computer systems, requiring fundamentally new approaches in order to provide the necessary security while operating at gigabit data rates.

In the perfect environment, the goal is to have computing systems capable of delivering data quickly and at very high speed among applications, peripherals and the network. While the work on network peripherals is clearly a job for networking researchers, much of this work must be done in collaboration with the operating systems and community and is therefore mentioned again under related research.

2.3. Gigabit Interfaces for PCs and Portables

As gigabit networking technology becomes pervasive, it should naturally propagate down to the mass-market workstation (PCs). These PCs will soon have local backplane bandwidth in excess of 1.2 Gb/s however they are typically ill-designed or programmed to handle large amounts of networked data. We need to encourage recognition of the need for network access to PCs at high speeds and for architectural support for networking in these PCs.

Hand-in-hand with developing PC architectures we should begin to consider the prospects of gigabit networking the portable PCs. Today's portables are typically only a year or two behind desktop PCs in performance and so could be gigabit capable within five years. Simply making gigabit access to portable computers a reality (regardless of whether via a fiber or wireless interface) requires attention to developing low power network interfaces which require minimal cooling and are small enough to fit into the tight space within a portable.

2.3.1. End to End Protocol Issues

Distributed system application performance in gigabit networking environments depends on the efficient and timely processing of the communication protocols used by the application. With existing communication systems, application protocols are assembled to meet an application's requirements by selecting protocols at each layer of a protocol architecture. Message encapsulation combines the layer protocols into a complete protocol that supports the application. The layered approach to protocol design produces correct protocols, but protocol performance may be suboptimal for several reasons. One problem is that the small number of protocols to choose among at each layer means that the application must often accept unneeded functionality. Example of functions that are often provided unnecessarily include sequencing, reliability, inappropriate buffer management, and unnecessary synchronization. Another problem is that some functions, such as multiplexing, occur in multiple layers, and thus, the same function must be performed multiple times. Ideally, an application should be able to select functions within a layer or among layers individually.

Today, choosing the specific protocol that is used to implement a function can be difficult, because each protocol is optimized for a different communication environment. Further, because networks will support quality of service, matching protocols to application requirement and network service is important for high performance. Because of the wide range of applications that will need to be supported, we expect that new protocols will need to be designed and evaluated and methods for dynamically combining protocol functions will need to be developed. We need to develop a systematic approach to protocol evaluation and design that takes both correctness and performance into account. Ideally, we would like to derive all families of protocols that correctly and efficiently implement a given function, and determine the performance relationship among these protocols.

2.3.2. Innovative Shared Media Access Technologies

Historically, researchers believed that gigabit rates would make it harder to develop new shared media access technologies. However, new research suggests that there may be a wide diversity of media access techniques that work at gigabit rates and that need exploration.

The design of a shared media access protocol tries to balance two requirements. First, the access protocol must be relatively fair. If multiple systems are competing for access to a link, if all systems have equal priority, they should all receive about equal shares of the link bandwidth. Second, access to the network needs to be timely, consistent with the requirements of the applications using the network, which has often forced relatively short limits on packet sizes due to serialization delays.

In gigabit networks, the bandwidth of the network is so large that the timely access problem largely goes away. The major problem is fairness, because fairness requires coordination between senders and the delay between senders (measured in bits sent on a link) is very long. However there has been new work in fairness protocols (like S++ and the DQDB access protocol) which which has shown that distributed algorithms can avoid the effects of delay while giving reasonably fair access. These results strongly suggest that there is a diversity of media-access techniques available for gigabit networks which, to date, have barely been examined.

2.3.3. Parallel Channels and Striping

Price-performance ratios often make it less expensive to provide high performance by aggregating multiple channels rather than deploying a single higher-bandwidth channel. Aggregation of multiple channels, or inverse multiplexing, is commonly referred to as striping. Striping can be used to achieve desired bandwidths using available channels, such as building an OC-12 equivalent from four OC-3c channels, as was done, for example, in the Aurora gigabit testbed. In some cases, striping is achieved using only the host interfaces, i.e. the network is essentially unaware that striping is being performed.

There are at least three distinct approaches to striping:

- Packet striping. Each packet is sent on one channel, but a host has access to *n* channels.
- Cell striping: One ATM cell is sent on a single channel, so that each packet is spread over *n* channels.
- Byte striping: Each cell or packet is distributed across the *n* channels.

Note that different channels may experience slightly or wildly different delays due to differences in path length, as well as queueing and multiplexing effects. This phenomenon is often referred to as skew. Each approach has both advantages and disadvantages. In the case of packet striping, the differential delay is taken care of by host, just as for IP packets (a plus). However, the latency experienced by any single packet is the same as for the unstriped network. In the cell striping case, for any packet at least *n* cells long, latency is 1/*n* of the unstriped case. However, differential delay substantially complicates reassembly. With byte striping, reconstruction of cells or packets must take place at every switch, which adds cost to the network, not just the high end users.

Several issues need investigation. The merits of the various approaches for a range of realistic applications need to be characterized. Accurate methods to measure and estimate the distributions of skew arising from various sources need to be developed. The potential limitations of striping, e.g., due to limits of synchronization schemes in the face of skew, buffer requirements, and errors need to be more fully explored. The cost benefit tradeoffs for striping of various flavors and at various values of *n* need to be understood. And finally, experiments with various promising striping mechanisms, using hardware, software or a mix of both, need to be conducted.

2.3.4. Design and Verification of Protocols

The desire to allow for modular protocol design, and also the desire to integrate multiple services on the same transmission technology makes systematic design and verification of protocols more important.

Protocols define the fundamental rules of operation for communication networks and are hence, central to ensuring that networks operate reliably. Logical correctness is one part of achieving reliable operation, but real protocols must also provide a level of performance that allows them to tolerate errors on the part of their environment. These errors may be caused by noise, hardware failures, misuse (either accidental or deliberate) on the part of end users or network operations staff, logical errors in other components of the overall network system or unforeseen interactions among different components. General methods for protocol analysis and design approaches that lead to inherently greater reliability are consequently of primary interest to network researchers. Society's growing dependence on communication networks makes it imperative that networking researchers develop systematic methods for ensuring that communication protocols exhibit both logical correctness and robustness of operation.

To date the field of protocol verification has struggled because of the complexity of completely analyzing the possible failure modes of communications protocols. Indeed it is fair to say that verification lags at least 15 years behind the data communications field as a whole, in that it is widely accepted that today's standard verification techniques are hard pressed to verify the reliability of protocols designed in the 1970s. Yet there have been recent advances in techniques that suggest our ability to design and verify protocols may improve.

One such technique is systematic design, where the design process takes a simple and verified design and through a series of well-defined transformations builds the more complex protocol required for today's networks. Because the transformations do not affect the correctness of the protocol, the resulting protocol is safe derivation of the original verified protocol and does not require reverification. In any case, the need for better techniques is clear.

3. Gigabit Applications Research

Gigabit networks are in their infancy; gigabit-speed links have been operational for a short time in only a handful of small network testbeds. Consequently, there is only a tiny base of experience with the use of this new technology to provide new services, solve previously intractable problems, or achieve higher productivity. Nevertheless, we can already identify a number of applications that require or would be greatly enhanced by the availability of gigabit network connections. Furthermore, when other technologies have had orders of magnitude advances in a short period, new, unexpected applications of the technology emerged. The two to three orders of magnitude increase in speed of networks that is now possible will surely spawn applications that have not yet been identified.

While it is impossible to predict the applications that will eventually pervade gigabit networks, we can begin to identify some types that will dominate for the next few years. They can be grouped in the following taxonomy, which is very similar to that presented in section 1.2.

1. Applications that are unique to gigabit environments, i.e., require gigabits speeds to be done at all (gigabit-enabled).
2. Applications which need to be radically revised or changed to perform well at gigabit speeds (gigabit-challenged).
3. Applications that have been already identified that scale to gigabits but only become practical, cost-effective, or pervasive as a result of gigabit availability (gigabit-enhanced).

The major difference between this taxonomy and that of section 1.2 is that the third class, gigabit-enhanced applications, is presumably within the scope of a gigabit research program because, while the application could be made to work at lower speeds, the resulting application is only of interest (and thus likely to developed) when gigabit networks are available.

3.1. Existing Gigabit Applications

A few applications have been implemented in the existing gigabit network testbeds and have shown the utility of such networks. Some Grand Challenge class scientific applications (e.g., quantum chemical reaction dynamics, global climate modeling, chemical flow-sheeting, traveling salesperson problems) have solved problems that were previously out of reach by using the network to combine geographically distributed computing resources. When heterogeneous computer architectures were involved, some applications achieved remarkable superlinear speed-ups. With slower networks, these applications would not have achieved significant (or any) speed-ups and thus the geographically distributed computing approach would have had no effect. Other applications interactively combined observational, possibly real-time, data (from a variety of sensors or instruments) with output from a computer simulation model to permit exploration of complex phenomena. Examples of this type of application include radiation treatment planning for cancer patients and predictions of location and magnitude of earthquakes. The geographic separation of data bases, instruments, supercomputers, and users plus the data-intensive and interactive nature of the applications mandate the use gigabit-speed networks.

Most of the existing applications involved transfer of data that could not be compressed if the compression resulted in any loss of accuracy. This was due to the fact that the data was being sent from one program unit to another for further computations. Loss of accuracy was unacceptable in these cases, whereas if only visual data to be seen by people were being transferred, some loss of data or accuracy might well have no harmful effects.

Parallel computing techniques were used for programming the metacomputer created by combining systems across the gigabit networks. Distribution of the application programs across multiple machines was typically based on functional decomposition. Different tasks or phases of the program were parceled out to the network-connected computers and these tasks exchanged information during the computation as needed. Pipelining approaches were used to hide the latency of network communications.

Some applications required the use of independently developed programs that were interfaced using the network. Thus programs that might be difficult to re-host could be run in their customary (but different) environments yet exchange data dynamically and thus improve the fidelity of the simulation by taking into account additional features of the phenomenon under study. This approach was also interesting as an experiment in facilitating collaboration among groups that were previously independent and in enabling the comparison of competing methods.

3.2. Predicted Gigabit Applications

The workshop found it useful as part of its discussions to try to identify typical gigabit applications of the future.

3.2.1. Telepresence

A class of applications that is on the horizon can be labeled as telepresence. Telepresence applications include teleconferencing, telemedicine, teleoperation, telescience, telecommuting, and telecollaboration. They all use multiple media streams with widely ranging properties, including high bandwidth real-time video (both high- and low-resolution), shared workspaces (text editors, drawing tools, and application-specific programs), and interaction control (floor control, user activity indication, invite/add/delete/join, etc.).

Data collection from high-data rate instruments can require substantial fractions of a gigabit/second. For example, in astrophysics, conducting searches for millisecond pulsars requires data collection at the rate of 400 Mb/s. Other instruments such as Synthetic Aperture Radar (SAR) can gather data at speeds in excess of a gigabit.

Perhaps the most compelling of the telepresence applications are in the field of medicine. Telemedicine can encompass a wide variety of tasks and applications. We list a few that have been extensively discussed as prime candidates for high-speed networking applications:

- Sending images and data from medical instruments (MRI, X-ray, sonograms) to experts for consultation. In this case, high-fidelity and rapid displays of the images are the chief requirement. There is no consensus yet as to whether lossy data compression would be acceptable, but many people feel that at least initially data transmission must be lossless. Some images, e.g., standard X-rays, are relatively small (a few megabytes) and could be sent quickly enough with networks of OC-3 speed. At present X-rays are typically sent by courier from one physician to another, with transfer times of a few hours to a day or more and with as many as 20% of the X-rays being lost in the process.
- Combining data from several instruments for the same patient to gain a more complete understanding of the patient’s condition. In this application again data such as MRI, X-ray, sonograms would be transferred, but in this case for the purpose of combining the different data sources with the aid of computers and then visualizing and navigating interactively through the fused data images. By seeing different instrument data simultaneously, physicians might get

additional clues of the nature of the problem. By traversing the visualized data interactively and experimenting real-time with opacity levels for the different types of data, they might discover clearer ways to identify the problem. Computers would be used for real-time data processing for image enhancement as well.

- A variant of the scenario above in which powerful computers are used to perform simulations (e.g., radiation dose computations, fluid dynamics) or to carry out a data base search to find similar cases. In the latter example, feature-based queries of terabytes of images might be required.
- Consultation via video teleconferencing, during which historical images as well as new images would be called up and viewed simultaneously by the participants in the videoconference.

There are many other aspects of telemedicine, but those sketched above suffice to identify the salient features.

- The application is of profound social and economic importance. Health care could be improved substantially by more frequent use of the top experts for consultation and by more efficient use of physicians' time.
- The experts, patients, hospitals, and data sources are very often separated geographically.
- Lossless transmission will typically be needed. This requirement has implications for the bandwidth needed (because the benefits of compression are limited) as well as for protocols that ensure delivery of the data.
- Medical data are very sensitive. Most if not all medical applications that require networks (at any speed) cannot be undertaken until the networks are demonstrably secure.
- The transmission speed required is not always in the gigabit range but the more interactive applications are expected to need gigabits per second to be accepted by physicians.
- Since some medical applications might be for life-critical services while others might be for background research, it would be very desirable to be able to specify different levels of service for different network applications.

Broadly speaking, most telepresence applications are gigabit-enabled. One can develop prototypes now, but to achieve the qualities required by users like physicians requires gigabit bandwidths.

3.2.2. Distributed Manufacturing

Distributed manufacturing is a second class of applications that are enabled by gigabit networks. It merges the telecollaboration environments of telepresence with remote simulation of components, both mechanical and electrical, and with real-time control aspects. Integrated design and manufacturing that involves many parts suppliers would present special challenges in the protection and interfacing of proprietary data and software.

The design and manufacturing of many products requires extensive simulations carried out on powerful computers and using sophisticated programs that have been developed over long periods of time. In some cases, remote data bases of parts, properties of materials (such as tensile strength), or instrument data must be consulted in the course of the design.

This application shares with telemedicine the requirements for protection of the data from loss during transmission as well as from being obtained by unauthorized parties. For many companies the computer programs that are used for design and analysis of their products are among their most precious assets.

Large manufacturing firms often have geographically-distributed engineering groups. Alternatively, companies use contractors to supply parts of the overall product. Sharing of engineering drawings and CAD/CAM data over networks can require high aggregate data rates. When dozens of engineers at one site are examining and browsing through detailed digital engineering drawings that are located at another site, gigabits flows are needed.

In some cases, it will be necessary to use networks to link proprietary simulation programs each of which addresses different aspects of the analysis or design of a product. By running the proprietary programs at the site of their owners, and having the programs exchange interface data across a network, one could avoid the concerns about intellectual property while gaining the ability to carry out global optimizations that are not possible today. In addition, there would be increased confidence that when the various components of the product are built they will interoperate as intended.

In summary, distributed manufacturing is gigabit-enabled because the high bandwidths (provided they are secure) would permit industry to bring to access sophisticated simulation programs, powerful computer facilities, remote data bases, and dispersed enginering groups and bring them all to bear on the design and manufacturing of products.

3.2.3. Distributed Computing

Distributed computing on clusters of workstations connected by LANs has quickly become a popular way for some applications to achieve quicker turnaround or tackle larger problems. Latency and bandwidth limit the classes of applications that can benefit from such an approach. Gigabit networks will broaden the spectrum of applications that benefit from distributed large-scale computing in the large by increasing the bandwidth of the connections and the grain size (computing speed and memory size) of the computing resources. The larger bandwidth can be used to at least partially offset the latency effects. The NSF Metacenter is an example of the computing environment that will be created by gigabit networks.

Many distributed computing applications are both gigabit-enabled or gigabit-challenged as they attempt to balance the greater available bandwidth with the problems of fixed delay. In many fields of science and engineering, mathematical models have been developed for phenomena to be simulated with great fidelity using computers. The pacing resource for doing bigger or better simulations is the computational power that can be applied to the problem. (By ‘‘computational power’’ we mean primarily computing speed, memory size, and IO speed). Since networks can be used to link large numbers of computers, a possible way to gain access to vast amounts of computing power is to distribute a computation across many computers located around the country. To do so, one must split the application into subtasks and run the tasks on separate computers.

In general, the pieces into which the application has been split must communicate with each other during execution. In wide-area networks, latency will perforce be high because of propagation time if nothing else. Since the speed of light cannot be improved, one cannot reduce propagation-induced latency. However, if exchange of data occurs infrequently, the computation can proceed efficiently despite high latency.

If there are large volumes of data to be communicated among the computers that are taking part in the computation, the distributed application will only proceed efficiently if the bandwidth is sufficiently high. Distributed scientific applications exist today that, when run on computers capable of gigaFLOPS require gigabit bandwidth to be efficient and, when run on the teraFLOPS computers that will become available in the next few years, require multi-gigabit bandwidth to be efficient. Such applications are examples of gigabit-enabled applications. When networks become faster, these applications will be able to tap efficiently much larger computing resources than is now possible and phenomena will be studied that are currently out of reach of computer simulations.

3.2.4. Distributed Information Services

New distributed information services will emerge with the availability of gigabit networks and large, distributed databases. The databases may consist of digital libraries, sensor data, or image data. The use of Mosaic to access the contents of the World-Wide Web (WWW) may be thought of as a precursor to the distributed information systems of the future.

It may well be that this class of applications will be the strongest consumer of gigabit networks. Large databases are harder to replicate than large computing systems. Many databases are too large or too frequently updated to be replicated at user sites. Furthermore, many predicted applications will use simultaneously data from multiple databases. Interactive analyses will preclude straightforward prefetching of data.

These services are likely to have the greatest impact as National Challenge-type applications, even though the WWW offers only pedestrian services. Distributed interactions between clients and servers can be combined with Grand Challenge distributed computation, such as remote sensing and integrated database systems. These systems can exhibit emergent properties more powerful than their individual components, by integrating the components of National Challenge applications.

Distributed information systems are interactions between information servers and information consuming systems over a distance. They differ from telepresence in that the server is typically automated, i.e., this is human-computer telepresence, whereas telepresence may include a human-human component. Human-computer interaction is constrainted by response time, usually on the order of 100 ms or so. Given speed-of-light propagation latencies of 30 ms (U.S.A., coast-to-coast) and including switching latencies of 20-30 ms (expected even for ATM), this leaves about 40 ms for the entire client server interaction.

Packaging of requests, transmission of requests, and service of requests, are together likely to account for another 5-10 ms. This leaves 30 ms for transmission of the entire response. Even for the current typical WWW file size of 40-100 KBytes, acceptable response time dictates a bandwidth of around 20 Mb/s to the user. With ever increasing file sizes, it won't be long before near-gigabit speeds will be required to support real-time distributed interactions with information servers. For example, animations of simulations of the Shoemaker-Levy 9 comet impact on Jupiter that were put on the Web were a few megabytes each and the authors would have made them much bigger if adequate internet bandwidth had been available.

Another case to consider is that of server interactions with a very large propagation delay (100-200 ms) or one in which the bandwidth-delay product is much larger than the size of the typical response, such as a satellite network. In either case, there is excess bandwidth but latency is a constraint. There are tradeoffs between bandwidth and latency that can be used to increase the effective throughput and decrease the response time of a distributed interaction, specifically where the interaction is hypermedia navigation (e.g., WWW).

As they are currently used, these these distributed information services would be gigabit-enhanced, because the quality of services such as real-time video, get better with increased bandwidth. The evolving use of these services as interactive interfaces is gigabit-enabled, because such use is not possible at current bandwidths.

3.2.5. Gigabit Applications Research Issues

From the research work to date, it has become clear that there are several research issues in gigabit applications that need to be investigated. In some cases these issues are obstacles to implementing certain classes of applications. In other cases, the issues represent impediments to making gigabit applications more widely available. The major issues are:

1. Many applications can be tuned to run at a gigabit, but the level of effort is huge, because we lack any general tools for gigabit application development. As a result, achieving gigabit performance is an application-specific affair. One of the problems with this approach is that, depending on how it is tuned, the application may have to be retuned as even faster networks are developed.
2. Even if we have gigabit links, applications can't necessarily use them effectively when the bandwidth-delay product is high. This is the case when the applications don't have enough data to send to keep the pipe full, i.e., they cannot take advantage of the available bandwidth.
3. Some kinds of applications cannot be written at all, including highly-networked applications that are dynamically self-adapting to a changing network and resource environment.
4. Proprietary or sensitive (e.g., medical) information and software needs to be protected. While such protection is needed even for low speed networks, gigabit networks present special challenges both due to the higher speed requirements for tasks such as encryption and decryption and the huge number of processes that may be involved.

The next section examines particular application research topics in need of attention.

3.3. Application Research Topics

Specific research topics motivated by the issues above can be split into two categories: key enablers and capabilities needed to make the gigabit applications practical or commercializable. Much of the research initially needs to be in ''middleware'' that the applications can use as building blocks or use for obtaining services or information. This focus represents the development of gigabit application expertise over the past six years. The problem is shifting from ‘‘can gigabit applications be written?’’ to ‘‘how to make gigabit applications more flexible and easier to write?’’

3.3.1. Programming Paradigms

The development of the remote procedure call (RPC) programming paradigm tremendously reduced the difficulty of writing distributed applications. Programmers were given a simple interface to the network which worked well enough over a variety of conditions. Unfortunately, classical RPC does not perform well on gigabit networks and there is a need for a new type of programming paradigm that allows programmers to easily utilize at least most of the benefits of using a gigabit network.

Several approaches have emerged as possible programming paradigms in gigabit networking environments. Distributed shared memory could present a consistent and familiar interface. Enhanced versions of RPC, allowing concurrent RPC or exportation of computation, could allow RPC to extend to gigabit speeds. And there is hope that object-oriented systems might use object migration to reduce latencies in response to changing access patterns. However, none of these technologies has left the lab and been made widely available to the applications community at large. In some cases there remain questions about the ability of the individual paradigms to exploit gigabit network performance while maintaining a simple interface that programmers find comfortable. Yet if gigabit applications are to become easier to write, an aggressive program to develop and disseminate at least some of these paradigms is needed.

3.3.2. External Data Representations

It is generally considered critical to heterogeneous distributed computing to have a common external representation for the transmission of a given data type, and make it the responsibility of the sender to convert its internal representation of the data into the external form, and expect the receiver to convert from the external form into the receiver’s internal form. There are several advantages to external data representations. They ease communication among heterogeneous system (a system need not know the data formats of others systems to communicate with them), and they can help minimize transmissions when multicasting to heterogeneous receivers (only one copy of the data needs to be sent, regardless of how many different data formats the receivers use).

However it is also true that the memory (and sometimes computationally) intensive conversion to and from external data formats is often the major bottleneck in deliverying gigabit throughput from the network to the application. At this point it appears there is no good alternative to external data representations, but it is true that some representations are less computationally expensive than others and some thought should be given to finding data formats best suited to scaling to higher speeds.

3.3.3. Distributed Control Mechanisms

Recent experience has shown that pipelined applications and applications that engage in coarse-grained distributed computation can both make very effective uses of gigabit networks. However we currently lack common tools for developing and managing such applications. In effect, each application programmer is forced to develop her own solution to the problem of controlling the distributed application. There is a need for common tools and programming interfaces, or at least guidelines, about how to write effective distributed applications.

3.3.4. Adaptive Applications

One of the exciting ideas in the past several years has been the notion of adaptive applications − applications that modify their communications behavior in response to the network service they experience. Adaptive applications can play an important role in gigabit networks. For instance, coarse-grained applications should be able to adapt the level of coordination between elements of the computation based on the experienced network delay. At a given level of performance, two systems in the same state should be able to exchange more frequent updates than two systems at opposite ends of the country. And to adapt to the longer delays, the two systems may wish to exchange larger pieces of information in the cross-country environment.

We need to better understand adaptive applications. Issues such as how widely applications can adapt and whether adaptations are necessarily application-specific or whether there are general rules that can be applied across multiple types of applications have yet to be explored.

4. Infrastructure in Support of Research

The workshop also considered the question of what research infrastructure is needed to foster and sustain gigabit research over the next few years.

The term ‘‘infrastructure’’ refers to metropolitan and wide area hardware, transmission, and software facilities in support of gigabit applications and networking research. Local area network (LAN) facilities are not included, nor is ‘‘middleware’’ that might be considered as infrastructure needed for applications research; the latter is assumed to be a either a networking or an applications research topic.

There is a continuous three-way interplay between application research, the desire to develop a user community which actually uses the applications, and network research, resulting in the evolution of the infrastructure. Applications researchers need an infrastructure on which they can develop and test applications. Network researchers use the infrastructure to develop and test hardware and software systems that will become part of the next version of the infrastructure. They also test their work with real traffic from the applications, especially traffic from applications run by real users, refine the systems accordingly, and incorporate the results into the infrastructure. The improved infrastructure can then be used to develop new or better applications, which in turn produce new traffic characteristics that influence the design of the next version of infrastructure. In summary, applications researchers gain insight from using the infrastructure while network researchers gain insight from building the infrastructure in response to how users employ the applications.

It is important for researchers to know that the infrastructure will persist. This longevity will encourage researchers to invest the time and effort to develop new capabilities. It is important to users both to know that the infrastructure will persist and that it will be relatively operationally stable. It is also important that the infrastructure become commercially viable and interoperate with other (i.e., lower-speed) elements of the National Information Infrastructure.

4.1. Scale

Scaling is a significant consideration to both network and applications research. Over the past several years, a handful of relatively small-scale gigabit testbeds have been investigating alternative architectures that might be used to construct a national-scale gigabit network. Although the first interplay between applications and network research occurred on these testbeds, the efforts were necessarily limited by the scope of the testbeds. For several reasons, the infrastructure for the next stage of gigabit research must be on a much wider scale. First, the insight gained with the existing testbeds demonstrated the utility of gigabit networks and hence there are more practical applications that should be developed. Second, the results of network research must be verified on larger-scale systems and must include the effects of traffic from multiple concurrent gigabit (as well as non-gigabit) applications. Third, it is important to provide testbed access to more researchers to enable them to experiment in a real network environment. Finally, it is important to begin educating the wider community in the uses and capabilities of gigabit networks, by making gigabit networking a part of their daily life. Fortunately, an increasing number of end-resources that are gigabit-ready could be utilized in a such a gigabit network environment.

4.1.1. Scale - Applications

There are two scaling issues in applications. First, one has to estimate the size of a network required to permit experimentation with new gigabit applications. Second, one must estimate the network size required to develop a significant user community that can provide useful feedback to researchers and form the core of a future gigabit networking user community on an NII.

4.1.1.1. Scale - Applications Research

It is projected that many new gigabit applications will require several sites be able to intercommunicate with multiple independent gigabit streams. The endpoints for these applications might be supercomputers, databases, instruments, virtual environments, and workstations. Therefore, to support, say, five such applications, the infrastructure must link approximately 25 to 40 sites, each with gigabit access (assuming some overlap of resource use).

An even larger number of sites may be required to support high-speed applications whose combined bandwidth needs require gigabit speed backbones. For instance, delivery of multimedia data may be done hundreds of sites simultaneously. Some distributed computing applications would benefit from high bandwidth access to several tens of computing sites.

As noted in earlier sections, the gigabit application streams mentioned above must flow all the way from end-to-end, from user-space to user-space. This means that in addition to the infrastructure itself, the host interfaces, the host operating systems, and the network APIs must support gigabit flows and preferably without dedicating the host to IO.

4.1.1.2. Scale - Applications Users

In order to have a large number of concurrent users of multiple applications, there must be a sufficiently large community of potential end users that we can reasonably expect them to concurrently use the network.

It is expected that independent of any government-sponsored program in gigabit network research, the number of end users with workstations or personal computers capable of communicating at speeds from 100 Mb/s to 1 Gb/s will grow substantially in the next few years, probably to tens or hundreds of thousands of users. The gigabit user community can presumably derived from this large pool. The challenge is to make sure the resulting pool of users is sufficiently large to provide enough diversity of application use, yet make sure each application has a large body of users who can provide feedback to application authors.

4.1.2. Scale - Networking

For network research, the scale of the infrastructure must be large enough to reveal routing and congestion issues and to develop quality of service schemes in a realistic environment. While initial tests can be done with lower-speed facilities and with a small population of users, proper verification requires gigabit speeds with a sufficient number of applications that are diverse enough to produce realistic congestion conditions and traffic/flow interactions. This is consistent with the needs of applications researchers for infrastructure that supports concurrent gigabit applications.

In order to uncover scaling issues, the scope of the infrastructure should be large enough to scale a given experiment up at least two orders of magnitude from early tests with a small configuration (e.g., five endpoints) to subsequent tests with, say 500 endpoints. The pool of a few thousand application users noted above would help ensure that the variety and quantity of traffic is adequate to verify the results of networking research.

There is an important tension in the infrastructure needs of networking and applications. Applications researchers and applications users need a stable network with enough bandwidth to effectively run their applications. Networking researchers periodically need to experiment with the network, which can make it appear unstable to the applications community.

To keep the network stable for applications researchers, some of the networking research can be done by isolating parts of a larger network for experimentation, for instance by connecting a few selected LANs on separate high-speed links. However, the networking community has learned by painful experience that there are few substitutes for experimentation with real, wide-area, networks. One networking aphorism states that local area solutions rarely scale to the wide area. So regular wide-area networking experimentation at gigabit rates is likely to be needed. This does not mean, however, that all networking researchers need high-speed connectivity. Some work can be conducted by researchers using relatively low-speed links to manage and observe experiments on the high-speed infrastructure.

4.2. Knobs and Levers

The tension between applications and network research can appear in issues of performance measurement and the plasticity of the networking infrastructure. Applications researchers typically desire predictable high-performance service in order to conduct their work. However, the problems of providing this high-performance service is a networking research issue, and research networks typically require networking components with extensive instrumentation (for measuring behavior) and considerable plasticity (so that the network component's behavior can be radically changed as measurements dictate). Unfortunately, it is difficult to build an infrastructure with sufficient plastic to support network research but which also works fast and reliably.

In general, we envision that the infrastructure will have to support at least three types of experimentation:

1. application experiments over links with a wide range of qualities of service (typically requested or specified by the applications)
2. network experiments that require specific traffic patterns, namely from traffic generators
3. network experiments that use application traffic

Observe that, if carefully partioned from each other, experiments in the first two cases should not interfere with each other, since both are attempts to isolate specific behavior. The third type of experiment requires that applications operate over experimental configurations and equipment that might result in network outages.

4.2.1. Knobs and Levers - Applications

The application community requires an infrastructure that is stable, though it need not be uninterruptible. The infrastructure must be available on a consistent enough basis that development and testing as well as scheduled demonstrations can take place. At the simplest level, the application researcher must be able to specify experimental conditions. For example, an application might need a very low probability of data retransmission. Similarly, the ability to degrade (or simulate degradation of) the network service is important for application experiments in order to develop and test robustness, adaptation, and fault tolerance schemes.

An application may also need information about the behavior of the infrastructure, including bandwidth, load, latency, etc., so that capabilities for determining resources, routes, and algorithms can be built into the application. For example, a distributed numerical model might require four types of resources, but there may be a half dozen of each type available on the infrastructure. The application would use information about available endpoint resources, network topology, and performance (including bandwidth, latency, current load, loss) to select among the possible permutations of these resources. Feedback from the network, in the appropriate form during execution is also desired so that the application can be designed to adjust to current conditions without restarting.

4.2.2. Knobs and Levers - Networking

The ability to manipulate multiplexers, switches, and other components of the infrastructure is essential for networking research. For example, researchers may wish to vary queueing disciplines within switches or to modify switches to include support for fast virtual circuit setup or multicast traffic. It is also important to provide the ability to collect statistics regarding traffic, loss, queue lengths, errors, and the like.

In addition, some network research might need access to the infrastructure's transmission facilities, e.g., SONET equipment or even dark fiber. For example, research in wavelength division multiplexing would benefit from wide area access to dark fiber. In addition, research using access speeds of 2.4 Gb/s (OC-48) and eventually 10 Gb/s (OC-192) would be desirable.

The technology used in the infrastructure should be diverse, in order to provide a testbed similar to that in the real world. Heterogeneity requirements dictate some diversity in the infrastructure and its interface such as use of raw ATM or SONET. A diversity of infrastructure interfaces should also be provided, including IP, HIPPI, raw ATM, and SONET.

4.3. Deployment Considerations

To meet the scale needed by the network research community, the infrastructure must eventually interconnect at least 1000 end systems distributed across a large geographic area. The large area is required to test scaling, latency, and complexity issues. The number of endpoints is required to achieve a sufficiently large user community. This section examines some of the practical issues in deploying a network of this size.

4.3.1. Interoperability Among Networks

The infrastructure must be interoperable with lower-speed infrastructure (e.g. the current Internet) at some base level of functionality, and there should be plans and mechanisms for incorporating the gigabit research infrastructure into operational networks. By making this a goal of the infrastructure, we will help ensure that the gigabit research properly feeds the development of the future NII.

It would be ideal to work with multiple carriers with route diversity (and provider diversity) among MANs and sites that are hundreds of miles apart. The provision of a national scale gigabit infrastructure is an excellent opportunity to foster more collaboration and interoperability among service providers. This would also facilitate future transition to commercial service.

4.3.2. Management Issues

To meet the different controllability needs of the applications and networking research communities a time-scheduled scheme or physically separable portions of the infrastructure could be implemented. An alternative approach would be to establish a virtual boundaries between research and applications infrastructure. For example, ATM virtual path mechanisms might, in principle, provide logically separate elements of the infrastructure that do not interact with each other.

The infrastructure need not be a single large system, and in fact there would be distinct advantages to having a set of interconnected subnets For example, a set of subnets would allow for experimentation with homogeneous architectures (within a subnet) as well as interoperation between architectures. However, the subnets must be interconnected at gigabit speeds because lower-speed interconnections would not support the end-to-end gigabit streams needed by applications.

There should be some level of central management and coordination of the infrastructure, particularly those portions which are globally shared. If the infrastructure is composed of subnets, then subnet components might be locally managed, but the interconnections among the subnets would be centrally managed. At times, a researcher may also want to gain control of the entire infrastructure or portions of the infrastructure, and there should be mechanisms to allow for this both technically and from a coordination and scheduling standpoint. This capability would impose the need for security for these mechanisms be substantial, and this is in itself might be a research issue.

Note that the central management style must be consistent with the needs of researchers. All network management information that the central manager collects should routinely be archived and made easily available to researchers (probably through an on-line archive). The research community should be able to direct the central manager to change the amounts and types of information collected, to better support research work. And the central manager should be aware that information about network outages (however much they may distress users) should be widely circulated, as the outages may signal unexpected problems that require research.

4.3.3. The Development of Gigabit LANs and MANs

It seems reasonable to assume that a national research infrastructure can make use of the cost efficiencies of LANs and MANs by promoting their development and attachment to a wide-area gigabit infrastructure.

The infrastructure for gigabit LANs is becoming increasingly available as campuses make a practice of using single mode fiber whenever buildings have to be rewired.

The costs of installing and operating gigabit links within some metropolitan areas are often substantially less than the costs of providing gigabit wide-area links, due to the presence of alternative local fiber providers and the large number of available rights of ways (subway tunnels, train lines, lines of site for microwave transmissions) in many cities. Furthermore, a single MAN connecting a handful of high technology sites can provide gigabit access to tens or hundreds of end systems.

Another benefit of LANs and MANs is that they can provide small scale research environments in which some experiments can receive initial testing before being deployed on the wide-area infrastructure.

4.3.4. Site Selection

The selection of sites should be based on several factors, including the unique resources available at the site for gigabit applications and the number of endpoints and users that the site will support for gigabit research. In addition, network researchers who want to use test traffic from application sites will require hands-on network support from those sites during their experiments. Site requirements such as these are not unlike those established in the late-1980s for connections to the NSFNET backbone. Those connections were provided subject to two criteria: A connected site was required to submit plans for extending the network on its own campus and for organizing an effort to provide economical connectivity among regional sites. This approach ensured leverage from NSF funding, and equally important, it ensured that every connection to the infrastructure would provide a growing number of endpoints. Given that, at least within small areas, gigabit networks are beginning to become affordable, such a leveraging program seems like the most effective use of government funds.

4.4. Recommendations

Based on the preceding analysis, the workshop recommends the following approach to developing a national gigabit research infrastructure:

- Provide a large community of potential users to support research in gigabit applications
- Support both applications and networking research on given infrastructure by scheduling (or by appropriate isolation mechanisms)
- Implement a large-scale testbed that provides heavily instrumented prototyping environment that mixes lower-than-gigabit speeds with some high-end MANs or LANs (comparable to a much-expanded version of DARTnet)
- Make components of the infrastructure sufficiently plastic that network researchers can change or replace them as dictated by applications work.

The following table quantifies these recommendations with a proposed timetable:

	1995	1997
Sites with gigabit access	20	50
Gb/s-capable end-systems (at sites with gigabit access)	50	1000
~100 Mb/s-capable end-systems (at sites with gigabit access)	500	>2500
Sites with ~100 Mb/s access	100	

The sites should be selected based on available resources for gigabit applications, which should include on-site support (e.g. consulting, visitor facilities) for researchers not located at those sites. The sites should also be selected based on their existing local gigabit infrastructure (number of endpoints/users that can be included in the community) and their plans for both increasing the number of local gigabit endpoints and working with other sites in the region.

The infrastructure should be based on standard gigabit protocols and capable of carrying IP traffic. This would allow for maximal interoperability with existing gigabit MANs and LANs and with the lower-speed networks (campus 100 Mb/s networks and the Internet).

5. Related Research

Research progress in gigabit networking is intimately tied to research progress in other areas. The introduction mentioned ways in which other areas need the development of gigabit networks to help them make progress. Similarly, progress in gigabit networking depends, in part, on research progress on work in areas like transmission technology, operating systems, computer architectures and security. This section briefly outlines some research problems which will need to be addressed if gigabit networking programs are to be successful.

5.1. Low-Cost Transmission Technology

One of the major inhibitors to the deployment of high speed networks today is the high cost of optical-electronic components and the transmission components that perform media adaptation (e.g. line coding, synchronization). These components constitute a large fraction of the cost of ATM switches supporting 150 Mb/s ports, motivating twisted-pair alternatives for connections between wiring closets and workstations. At 600 Mb/s, the current cost of optical devices and SONET transmission components for one end of a link is over $2,500, giving a cost of $5,000 per connected workstation. This is about two orders of magnitude too high for wide-scale deployment, and as one goes to higher speeds the gap gets even worse. Indeed, it is currently less expensive to build high bandwidth links by aggregating a number of low bandwidth links than by using higher bandwidth components. While manufacturing economies of scale can be expected to bring component costs down, they are unlikely to close such a wide gap in a short time.

While research on such optical and transmission components lies outside the scope of what is generally considered networking research, these components have such an overwhelming impact on gigabit networking that research directed toward achieving fundamental improvements in the manufacturing economics of these devices would be of great value. (It should be noted that encouraging improved manufacturing technologies may be difficult as the major manufacturers are outside the United States). At the same time, alternative transmission technologies that can be implemented more cost-effectively should be explored. It may be, for example, that lower costs could be achieved in the end-system connections by exploiting their more limited distance requirements and the fact that they do not need the complex maintenance features built into transmission formats designed for long-haul facilities.

5.2. Faster Transmission Technology

If one of the goals of gigabit networking is to leapfrog networking technology forward to a new level of performance and then maintain that performance vis-a-vis other technology into the future, then we need worry about ensuring that the process is in place to continue delivering ever higher performance as other systems get faster. In at least one area, the problem of transmitting faster through a single connection point there is some concern we are not pushing performance swiftly enough.

Most of the research in networking at multi-gigabit rates is focussed on parallel access (many computers or users each having a one-gigabit attachment). However it is not unreasonable to assume that individual systems will want multi-gigabit attachments in the future. The U.S. Government has funded some innovative work in this area, but because of a shortage of industry research the aggregate research work being done on signaling at multi-gigabit rates may not be sufficient. For instance, considerable work still needs to be done on the next generation SONET transmission speeds of 9.6 Gb/s (OC-192) and 38.4 Gb/s, both of which are needed for high-speed speed ATM, even though some Government agencies have identified current needs for OC-192c technology.

5.3. Operating Systems Performance

Ultimately, the goal of gigabit networking is to allow applications quickly to accept from and deliver data to a network at gigabit per second rates. However, an operating system sits between most applications and the network, and one of the lessons of gigabit research to date is that the operating system is typically a bottleneck, even after the networking code in the system has been optimized. The network code must deal with other parts of the operating system such as memory and IO management subsystems, context management subsystems and schedulers, and these parts of most operating systems are typically hard pressed or unable to handle gigabit data streams, even if the underlying hardware is gigabit capable.

Although expertise exists in the networking community to tackle many of the operating system issues, it seems likely that operating systems would benefit from attention from a wider community. Many of the obstacles encountered when implementing high speed networking also affect developers of other high-speed peripherals such as the large high-speed storage systems needed to support video-on-demand and image retrieval. In other words, the desire for nimble operating systems capable of delivering large amounts of data in short amounts of time appears to be universal. Similarly, some of the hardest problems in networking, such as achieving high networking performance on multiprocessor computers, are problems for which outside experts may have valuable advice that the networking community needs.

5.4. Computer Architectures - IO Busses and Memory Subsystems

As computer networking leaps ahead in performance, it is becoming clear that some other parts of the computing environment need to be pushed to keep pace. In particular, the performance the IO busses used to connect peripherals to the processors and their memory, is often woefully inadequate. It is not uncommon in today's computing environment to build computer systems with near-gigabit network interfaces and processors capable of feeding those interfaces at high data rates, only to find that the bus between processor and interface is capable of only a few hundred megabits of throughput.

Memory subsystems also remain a problem. While innovative work has been done to improve the performance of memory chips, it remains true that slow main memories, or poorly designed interfaces between the memory subsystem and IO and processor subsystems is often a critical impediment to high performance.

It is important to keep in mind that these problems are not unique to gigabit networks but rather limit the performance of nearly all high performance peripherals.

5.5. Security

Classically we think of information security as being divided into the functional areas of access control, authentication, privacy, and data integrity. Gigabit networking performance for privacy, and data integrity is limited by the performance of the fundamental building blocks of security: secret key encryption, public key systems, and message digest techniques. It is possible, but less obvious that access control and authentication concerns are going to be affected by increasing network speeds. Experience has shown that security must be performance transparent to be widely used.

Encryption technology supporting gigabit rate communications is currently available only in the form of experimental devices. Widespread acceptance of a standard as a critical factor to commercial gigabit rate device availability. The probable de-certification in 1997 makes it unlikely that DES chips will be available at speeds faster than the 150-180 Mb/s range now on the market. Clipper/Capstone based devices may become available at gigabit rates, but given the unpopularity of the scheme and the recent decision not to encourage use of Clipper/Capstone for data, it seems unlikely it will be accepted widely enough to result in availability of gigabit rate capable devices. A further problem is that it is not clear that there exists any program to ensure that encryption techniques keep pace with advances in transmission speeds, offering the possibility that transmission techniques will rapidly outstrip our ability to send data over them securely.

The use of encryption on ATM networks leads to several new concerns. Research is currently underway with key agile encryptors at OC-3c and OC-12c rates. Relating key management to ATM call setup is a complex area. The current efforts in this area do not address multipoint to multipoint communications, which will be important for loosely coupled distributed supercomputing problems, (although this is not uniquely a gigabit networking issue).

Existing techniques for providing data integrity largely rely on software implementations and are targeted at traditional file movement. Integrating message digest and digital signature functions into the protocol stack may be desirable for providing real time data integrity and authenticity checks. Existing algorithms appear adequate for gigabit networking, but the implementations must be rethought to provide the required speed. Available implementations of such techniques as RSA and MD5 are one or more orders of magnitude too slow for gigabit networks, and this will remain true until encryption technologies are able to track transmission system performance more closely.

5.6. Programming Tools

Hand-in-hand with the shortage of tools to develop pipelined and coarse-grain distributed applications are tools to debug and validate them. Programmers on single machines can now make use of sophisticated source level debuggers with a rich set of debugging tools. There is no equivalent capability in distributed systems of the type being developed on gigabit networks and this lack is a considerable hindrance to the development of distributed gigabit applications. Much of this work needs to be done in concert with work on programming languages and environments in general.

Appendix A: Participants in the Meeting

This report was prepared by Craig Partridge (meeting chair) based on material provided by Bruce Davie, Roy Campbell, Charles Catlett, David Clark, David Feldmeier, Ray McFarland, Paul Messina, Ira Richer, Jonathan Smith, James Sterbenz, Jonathan Turner, David Tennenhouse and Joseph Touch, and minutes of the meeting. This report has been reviewed by the participants of the workshop, who are listed below.

Andres Albanese
Mehran Bagheri
Sujata Banerjee
Jeanette Barker
Scott Behnke
Richard Binder
Jevad Bouromand
Larry Bowman
Roy Campbell
Doug Carlson
Charlie Catlett
Bilal Chinoy
Dave Clark
Josh Criss
Bruce Davie
Steve Deering
Kai Eng
Vakil Faramak
David Feldmeier
Allan Fisher
Darleen Fisher
Andy Grimshaw
Fred Heaton
Raj Jain
Bill Johnston
Robert Kahn
Tim Kuhfuss
Barry Leiner
Victor Li
John Limb
Carl Love
Michael Luby
William Marcus
Ray McFarland
Gary Minden
Paul Mockapetris
John Morrison
Doug Niehaus
Alfred E. Nothaft
Hilarie Orman
Sean O'Malley
Dhabaleswar Panda
Craig Partridge
Ira Richer
Tim Salo
George Seweryniak
Allyson Showalter
Jeff Silber
Dave Sincoskie
Jonathan Smith
Mike Sobek
Stephen Squires
Saragur Srinidhi
Martha Steenstrup
James Sterbenz
Dan Stevenson
Michael StJohns
David Tennenhouse
John Toole
Joe Touch
Johnathan Turner
Walter Wiebe
Stephen Wolff
Yechiam Yemini
Lixia Zhang