Gigabit Networking Survey

Bobby Vandalore

This is a survey paper on gigabit networking.

Table of Contents

---

1. Gigabit Networking - Introduction

Computers and their attachments (like networks and disks) are getting faster everyday. The current CPU speeds of processor like DEC Alpha and Pentium are well over 100MHz allowing them to perform billion instructions per second (BIPS). This speed is comparable to supercomputer's speed five years ago. With the growing speed of computers the applications which run on them are now ranging from interactive graphics, voice recognition, video conferencing, real time animations etc. All these new applications will use networks to carry more data.

Network bandwidth is also increasing concurrently with the CPU speeds. When in 1980s 10Mbs Ethernet was considered fast, we now have 100 Mbs Ethernet. The bandwidth is approaching the speed on 1 billion bits per second (1 Gbs), much due to the research in the field of fiber optic signalling.

The three main fields data communications, computing and telecommunications are undergoing a period of transition. The field of computing is rapidly advancing with processor speed doubling ever year. The latest **RAID** (Redundant Arrays of Inexpensive Disks) has given rise to file-systems with gigabit-bandwidth.

The field of data communications which facilitates the exchange of data between computing systems has to keep up with the pace of the growing computing technologies. In the past the data communications provided services like the *e-mail*. Now applications like virtual reality, video conferencing, video on demand services are present.

For a century the telecommunication industry has been carrying voice traffic. This scenario is changing with telephone networks carrying more data each year. The data being carried by telephone network is growing at 20% per year compared to voice traffic which is growing only at 3% per year. Soon the data traffic will overtake the voice traffic. All this, has made the telecommunication industry more interested in carrying data in their networks.

So the three communities are now converging with common interests of carrying more data at higher speeds. This has led to some joint activities. The most notable of these activities is that which has led to the setting up of gigabit testbeds in United States. Other joint activites are the standardization of ATM (Asynchronous Transfer Mode), a suite of communication protocol to support integrated voice, video and data networks. Some organizations which are doing research in gigabit networking are National Coordination Office for HPCC (High Performance Computing and Communications), The Corporation for National Research Initiatives, IEEE Communications Society Technical Committee on Gigabit Networking.

When the gigabit networking was in its horizon, many researchers felt that the current knowledge about networking would not apply to gigabit networks which are considerably faster than existing networks. Now, after several years of research it has been found that many of the strategies and techniques (like layering the protocol) still work in gigabit networks also.

There are many working Gigabit testbeds (see AppendixA for full list). In five to ten years Gigabit networks will become a reality. It is now unclear whether there will be a single gigabit technology with a specific standard protocol. But it looks like that there will be many competing gigabit networking technologies (like many LAN technology) and many protocols but eventually one of them will become most popular (like IP).

The next section deals with the key concepts and technologies in the Gigabit networking. The third section deals with more specific issues in gigabit networking. The fourth sections discusses the various potential gigabit applications. The last section overviews the current state of gigabit networking. The appendix A gives a list of gigabit testbeds. The appendix B is an annotated bibiliography of the http sites, articles, papers and books refered in this paper.

Back to Table of Contents

2. Gigabit Technologies and Concepts

The development of high speed networks is closely linked with the advancements in fiber optics. The advent of fiber optic signalling equipment capable of transmitting at several gigabits per second over long distances and with low error rates through optical fiber showed that gigabit networks were feasible and has served as a goad to researchers. Other media like radio and micro waves also have been explored and found to be capable of providing gigabit bandwidth. The LuckyNet is a 2.4 Gbs (OC-48) testbed link build by AT&T. The Advanced Communications Technology Satellite (ACTS) is an experimental satellite that was launched by NASA in July 1993, which supports OC-12 links in one configuration.

Another important trend in gigabit networking is the increasing interest in a technology now widely known as *cell networking, cell switching or cell-relay*. In the following (sub) sections we will give a brief overview of these concepts and technologies.

2.1 Fiber Optics

2.1.1 Fiber Optic Basics

Some part of the light gets reflected when passing from of one medium to another, and the rest of it gets refracted. Light has an interesting property that if the angle of incidence is greater than a critical angle that all of the light is reflected. Fiber Optics uses this property of the light in sending the signals.

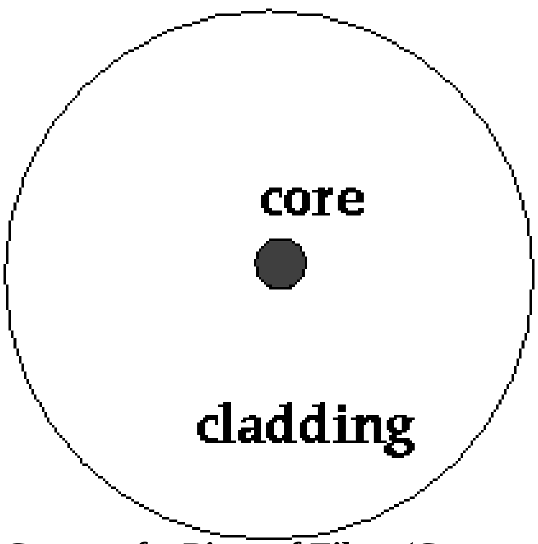

Center of a Piece of Fiber (Courtesy "Gigabit Networking", Craig Patridge)

The fiber has a thin strand of glass called *core* surrounded by a thicker outer layer *cladding*. Light is sent at the appropriate angle inside the core and it travels through the core, any light escaping the core will be reflected back into the core. The pulses of the light carry the bits. One important thing here is that the bits dont travel faster (propagation delay is similar to that of copper wire) than in copper wire, the higher bandwidth is got because the bits can be packed more densely (1000 times more than that of copper). Theoretically the fiber has bandwidth of 25 THz around each wavelengths of 0.85, 1.3, 1.5 microns. (These give rise to bands similar to those of radio waves). So the total capacity of a single fiber is 75 Terabits per second!

In fiber optics just like copper wire there are some problems while signalling. There are three major types of dispersions: modal, chromatic and material. Repeaters and Amplifiers which strengthen the signals are used to overcome these dispersions. Also *single-mode* fiber (also called *monomode*) is used to avoid dispersion. The *multimode* fiber still suffer from modal dispersion problems.

Transmitter and receivers are generic terms for devices attached to a fiber to respectively transmit and receive signals. These have two important varities: fixed and tunable. Fixed ones are set to a particular wavelength and tunable ones can dynamically set the lightwave frequency at which they transmit or receive. Fixed ones are simple, the tunable onces are more complicated. One example of a tunable device is the Mach-Zehnder Inferometer. Here one path of the light is made slightly longer than the other such that there are out of phase and the difference is used for tuning to a particular wavelength

2.1.2 SONET and the SDH

SONET the telephony standard which stands for Synchronous Optical Network (called the Sychronous Digitial Hierarchy, SDH in Europe). SONET was developed to support multiplexing on links capable of data rates of hundreds of megabits or more. Its main goal was to provide a single set of multiplexing standard for high-speed links. The standard provides with various rates called as *Synchoronous Transport Signal levels* (STS) or *Optical Carrier level*. The rates range from 51.84 Mbs (STS-1,OC-1) to 2.4 Gbs (OC-48). SONET carries the data in frames. Each frame has overhead and payload part.

2.1.3 WDM Networks

In contrast to the SONET, WDM networks uses the special properties of the optical fibers. In WDM networks the bandwidth of the fiber is divided into multiple channels and hosts communicate to each other on a particular channel. There are two major types of WDM networks: *single-hop*, *multihop*. As the name suggests in *single-hop* WDM networks the hosts are directly connected to each other via a star-coupler. The *single-hop* can be further classified depending on whether they use fixed/tunable transmitters and receivers. Two examples of *single-hop* WDM networks are LAMBDANET a project of Bell Communications Research and RAINBOW a project of IBM. The *multihop* WDM Networks can be designed in many ways. The main goal is to build a high-connectivity graph so that the number of hops between any two nodes is minimized. An example of *multihop* WDM network is the TeraNet, which was built as the part of ACRON project at Columbia University.

Back to Table of Contents

2.2 Cell Networking.

2.2.1 Fundamentals

The ideology of cell networking is that all the data should be transmitted in fixed size packets called *cells*. In networks usually the data is sent in *packets* which vary in size. If the packets vary in size it is difficult to guarantee bounded delays which are required for isochoronous traffic. See figure below

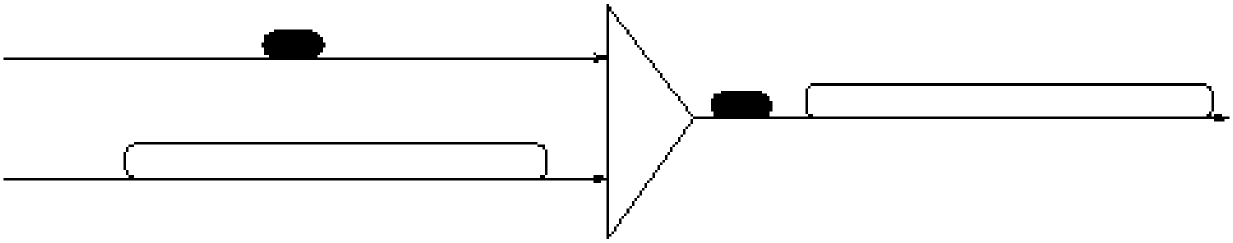

Serialization: Small Packet Caught Behind Big Packet

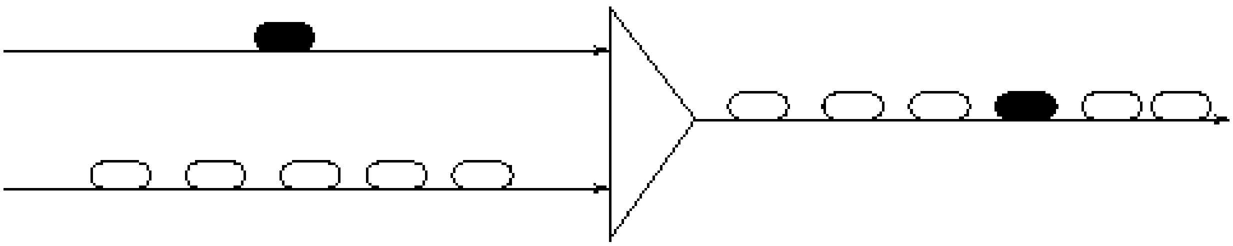

Serialization with Cells

(Courtesy Gigabit Networks, Craig Patridge).

The waiting time is large when the small packet of the isochoronous traffic (eg. voice traffic) is waiting behind a large packet. In contrast in cell networking, since the large packet is divided into small cells the waiting time is not large. In cell networking it possible to give guarantees for delay.

There are two ways of switching cells, one method is *store-and-forward* and the other is *cut-through*. In *store-and-forward* switching the whole cell is received and then it is forwarded in a appropriate link. In *cut-through* switching the switch decides on which link to forward the cell by just examining a few bytes of the header in the cell. Asychronous Transfer Mode (**ATM**) is the brain-child of the cell networking technology. The telecommunication community has put in a lot of effort to standardize this technology. There is also a *ATM Forum*, which is a group which is deciding on the main issues of the ATM before waiting for the official standard come out. Currently ATM provides bandwidth from 45 Mbs to 622 Mbs. But in future, ATM will also provide gigabits of bandwidth.

2.2.2 Cell Networking in LANs

Most of the Local area cell networking share the media. Since the media is shared the issues of who gets the access of the media has to be taken care. The Local area cell networks solve this problem by arbitrating the right to send a cell. Most the local area cell networks are ring networks. Now we shall see some examples of local area cell networks.

CBR

Cambridge Backbone Ring (CBR) was developed as collaborative project between University of Cambridge and Olivetti Research. It is a ring network in which the ring is divided into *frames*. Special five-bit patterns are inserted between the frames for synchronization and to equalize timing. CBR allows slots to be grouped.

IEEE 802.6 (DQDB)

The Distributed Queue Dual Bus is joint standard of Institute of Electrical and Electronic Engineers (IEEE) and the American National Standards Institute (ANSI). DQDB is the fore-runner of ATM, it has same cell format has the ATM. It uses a dual bus in a slotted-ring network. Each slot can hold one cell of 53 bytes (48 data + 5 header). The dual bus is advantageous because the network can reconfigure to form a bus when one node fails. Unlike the CBR, DQDB does not allow cells to be grouped.

HANGMAN

This is a prototype gigabit cell LAN built by Hewlett-Packard Laboratory in Briston, England. This differs from the CBR and DQDB in that it uses a large cells (256 bytes) because most of the LANs use large packets. It uses a logical folded bus architecture, each node is attached twice, one with the write bus and one on the read bus. One nice feature of the HANGMAN is that if the node wants to send a packet which is twice as long as the cell, then it can request for two consective slots and send the packet.

Another issue in local area cell networking is *cell-switching*. The **AN2** is ATM switch was built by Digital Equipment Corporation's Systems Research Center. It is a small (16x16) cross bar switch. The traffic is divided into continuous bit-rate and variable bit-rate virtual circuits. It can support data speeds as high as gigabit per second. The continuous bit-rate traffic must preallocate the bandwidth with the switch. To prevent loss in variable rate circuits, the switch uses a combination of input buffering and hop-by-hop flow control technique. It has a auto-learning feature, by which the switch learns the presence of other AN2 switches and this learning can be used for routing packets.

2.2.3 Wide Area Cell Networking

Switch design is key issue in wide area cell networking. The switchs will have to switch reliably very large bandwidths since it is operating at gigabit speeds. Usually these switches are designed using parallel interconnection devices. A fundamental problem in switching is blocking. *Blocking* happens when two cells content for a particular link. Another important issue is the buffer management. Various buffering strategies like *input-buffering*, *output-buffering*, *internal buffering* are used. We shall briefly see some of the wide area cell networking switches.

Crossbard Switches

In these switches every input is connected to every output by a cross bar. These switches use output buffering (cells are buffered at the output port). Multicasting can be easily provided in cross-bard switching. The crossbar switch usually have very low blocking probablities. Hence these switches are used as a standard for comparing other switches. Two examples of crossbar switches are the *knockout switch* and *Guass switch*.

The Banyan-Batcher Switch

The main disadvantage of the crossbar switch is that it uses O(n*n) circuitary for connecting n output ports. This switch uses only O(n*log(n)*log(n)) switches. This switch assumes that the n links have uncorrelated (independent) data links. It uses the batcher sorting element for providing connectivity from any input node to any output node. The *Starlite* and *Sunshine* are examples of Banyan-Batcher switches.

Back to Table of Contents

2.3 Packet Networking

Packet networking has been around since the advent of networking field. More research have been done on packet networking compared to cell networking. At gigabit speeds since the propagation delay of the medium is comparable to transmission time of the packet (eg 1.2 microseconds for transmitting 1518 bytes at 1 Gbs), problem due to serialization (large waiting time of small packets) are not present. Another problem the packet networks have to overcome is that they have to support isochoronous traffic. It has been shown that adaptive applications (applications which change their behaviour) can be used to overcome the problem of providing isochoronous traffic. Already some high-speed LANs and WANs have been built using the packet networking.

2.3.1 Local Area Packet Technologies

The ATOMIC LAN

This is a gigabit LAN built by USC Information Sciences Institute. It is built based on *Mosaic* chip developed by CalTech. The Mosaic chip uses a 4x4 or 8x8 mesh to interconnect 0.5 gigabit links. The chip routes packets by having *x* counter and *y* counter. The difference between the source and the destination is used to move the packet along the x and y coordinates. Mutlicasting is supported by having the processor and each concentartor to duplicate the packets.

The CSMA/RN Ring

This is a ring network (a follow on of the FDDI), developed and studied in Old Dominion University. Data from the ring is put in a short delay buffer. If the ring is not active (no data in it) then the controller can send the data immediately. Collisions are avoided by placing the new data from the ring at the end of the to-be-transmitted packet in the delay buffer. Another good feature about CSMA/RN is that the packets are removed at the receiver.

2.3.2 Wide Area Packet Technologies

The main problem of wide area packet networks are the router which have to route packets at gigabit rates. Forwarding can be done using very small number of instructions (100-200) instructions, because for forwarding, the router just examines the header and makes some consistency checks and routes the packet to the appropriate ports. Using multi-function instructions, the CRC calculation can be done while the packet is being copied itself. Also various improved hashing and lookup table techniques have been developed which are used in these routers. We shall see two high speed routers

plaNet - A High Speed Packet Router

It is a 8 port ring-shaped bus switch developed by Cidon and Gopal of IBM. It uses source routing, in which each adapter in the ring examines the packet and transmits the packet if the source routing specify it to be transmitted. Multicasting is provided by adapters replicating the packet. The packet is not removed after transmission, it circles round the ring once and then removed.

The Bell Labs IP Router

A team at AT&T Bell Laboratories has developed a prototype gigabit IP router. It uses multiple processors, the IP packet header is stripped and passed onto one of the processor which updates the header and sends it to the appropriate outbound processor. Its rate is just over a gigabit, but it is precursor of the future gigabit IP router which use multiple processor for routing.

Back to Table of Contents

---

3. Trends and Issues in Gigabit Networking

There is voluminous amount of research in the area of gigabit networking currently. The IEEE Communications Society Technical Committee on Gigabit Networking has been conducting a yearly Workshops on Gigabit Networking. In the following sections a brief summary of current trends and issues in gigabit networking is presented.

3.1 Challenges in Gigabit Networks

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, and to encourage the transistions of gigabit technologies into the wider community. To achieve these goals the Gigabit Netowrking Workshop identified important problems in the following areas :

Performance evaluation

> Higher speeds and new traffic mixes are causing a much needed re-examination of models and algorithms for networking performance.

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.

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

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 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.

Interfacing computers and application 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.

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.

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.

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 emering research suggests that there may be a wide diversity of media access tecniques that work for gigabit networks and these options should be explored.

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.

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 verfication technology is nearly 15 years behind the rest of the field, we need to encourage new work in this area.

Back to Table of Contents

3.2 Lessons from Gigabit Testbeds

Through a collaboration of industry, academia, and goveernment, work on five US testbeds was begun in 1990 with funding from National Science Foundation (NSF), the Advanced Research Projects Agency (ARPA), and industry. The testbeds are known as *Aurora*, *Blanca*, *Casa*, *Nectar*, and *Vistanet*.

The following are the lessons learnt from these and other Gigabit Testbed initiatives.

- The existing gigabit networks testbeds have proved that gigabit networking technology is feasible. So, the question is not "Whether we will have gigabit networks in future?", the question is "When we will have gigabit networks available?"
- Initially researchers believed that the existing networking techniques will not be applicable in high bandwidth. The testbeds have proved that most of the existing knowledge can be applied to gigabit networking as well. One example is layering of the network works well in gigabit networking also. There have been TCP/IP implemenations at gigabit speeds in the gigabit testbeds.
- A few applications have been implemented in the existing gigabit testbeds and have shown the utility of such networks. Some Grand Challenge class scientific applications (e.g quantum chemical reaction dynamic, 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 architecture were involved, some applications achieved remarkable superlinear speed-ups.
- Parallel computing techniques were used for programming the metacomputer created by combining sytems across the gigabit networks. Distribution of the application programs across multiple machines was typically based on the functional decompostion. Different tasks or phases of the program were parceled out to the network-connected computers and these s or pexchanged information during the

computation as needed. Pipelining approaches were used to hide the latency of networks communications.

- Gigabit testbeds have also helped in indentifying the networking problems which arise due to high bandwidth. Many of these challenges are mentioned in the previous section.

Back to Table of Contents

---

4. Gigabit Applications

Most of today's data network applications are not very sensitive to delay and variations in bandwidth. It does not matter very much if your files take a little longer to travel across the net. But in telecommunications (telephone) industry the applications are delay sensitive. Normally, when humans speak they pause in between sentences and if the pause is longer the other speaker speaks. If due to network delay if the pause is large then both the speaker may speak at the same time leading to confusion. So delay should be bounded in telephones. Other applications like X Windows, remote login etc will be faster if gigabit networks are used.

Any applications which needs low response time or high bandwidth is suitable to be a gigabit application. The recent advent of gigabit networks have given rise to many new applications. One of them is IVOD (Interactive Video on Demand). Here the consumers order which ever program they want to see and the programs are sent from a centeral server to the consumers. Since video applications use a large bandwidth and also different viewers may want to see different programs at the same time, this application will benefit from gigabit networking. Though compression methods like *MPEG* may be used for compression, every once in a while full screen data has to be sent. This can be done using gigabit networking technology.

Highly computation intensive problems can be broken into smaller problems and given to computers with high bandwidth networks connecting them for interchanging data. For example in UCLA, researchers are experimenting with simulation studies of atmosphere and ocean interactions. One supercomputer (CM-2) simulates the ocean and another simulates the atmosphere and these interchange huge amounts of data and the interactions are studied. Typically 5 to 10 Mbs of data is exchanged per cycle, this will take a second in a 10 Mbs Ethernet while it will take only 100ms in a gigabit networking environment.

Another class of applications are those which have real-time interactions with humans. A typical example is video conferencing. Humans are capable of absorbing large amounts of visual data and are very sensitive to the quality of the visual data. Another class is the *virtual reality* applications which give the user the illusion of being somewhere else. There have been interesting experiments done in NASA. They developed a system, by which the geologists can interact with the surface of the Mars. Geologists study by interacting with the surface, touching it (virtually), seeing the 3D scene from different angles etc. All these require large amount bandwidth, and gigabit networking comes to help them out.

One of the main difference between the traditional data-communication applications and those of interactive nature is that the later need timing requirements about spacing between samples. Recently there has been some work in some innovative experiments with *adaptive applications*. These applications change their behaviour dynamically and require only loose performance guarantees from the network. An example of the adaptive application is the *vat* voice-conferencing system developed by Van Jacobson. *vat* is like a telephone but uses the computer and internet to connect two peoples. The *vat* avoids isochoronous samples by keeping a large buffer and timestamping all the data it receives. Traffic which arrives earlier are buffered appropriately and then played. The time of the delay is called the *playback* point. So adaptive applications could have a major impact in network designing in future. The vat also has demonstrated that slow networks like Internet can support real-time applications with enough buffering.

Back to Table of Contents

Summary - State of Gigabit Networking

The first steps towards gigabit networking have been taken by exploring the capabilities of various media which can support gigabits of bandwidth. There is considerable amount of research being done in various gigabit testbeds, but still gigabit as field has a long way to go before becoming mature. One of the main issues is what protocol should gigabit networks run on. It is interesting to note here that, traditional protocols like IP, TCP with various modifications and extensions have been shown to support gigabit bandwidth. But should we stick with these are not is a question yet to be answered. Most of these protocols have not been verified because they have a huge number of states. Protocol verification is an important issue which has to answered by gigabit networking since at such high speed something can go wrong very fast. Some of still unexplored issues in gigabit networking is Network management. It is unclear whether the existing network management will still work at gigabit speeds. Encryptions of data takes some time, so fast encryptions algorithms (like DES, Digitial Encryption Standard) should be developed if authentication is to be supported by gigabit networks.

In most of the operating systems since the network interface is not well integrated, there is lot of redundant copying of data. So operating systems face the challenge of changing to support better network interfaces. This will become critical in supporting real-time systems. Some work has already been done in this area with improvements in processor speeds, better scheduling and caching techniques.

Traffic modeling is another area where work has to done. It has been recently shown that Ethernet traffic is Self-Similar (fractal) in nature. It has to verified whether this hold for gigabit networks or come up with new models for traffic analysis.

Will the network world stop progressing after gigabit networks or will terabits network be developed ? It is already been known that the potential of a single fiber is in the order of terabits. We can envisage terabit speed networks were a parallel array of high speed computers are connected. But will there be terabit applications to run on it? As it happens most of the time an application will be found whenever the technology has been invented. So we can expect to be working with terabit networks in the future.

Back to Table of Contents

Appendix A - Gigabit testbeds.

ACORN

Aurora

- MIT TNS WWW Site

BAGNET

- Bay Area Gigabit Testbed (BAGNET)

BATMAN

- BATMAN

BERKOM

- German National Research Center for Computer Science,GMD

Blanca

- UIUC NCSA Blanca Site

California Research and Education Network (CalREN)

CANARIE National Test Network

- CANARIE- " World's Largest ATM Network"(?)

DFN- German Research Network

- DFN

European ATM Pilot Network

- **European ATM Pilot Network**

Exploit Testbed

- Exploit testbed

MAGIC

- University of Kansas-TISL Info on Magic

MultiG, Sweden

- Center for Computing & Communication

Nero

- Nero

SuperJANET

- JANET News
- SuperJANET ATM Technical Advisory Group - document store
- SuperJANET Lower Layers Technical Advisory Group
- SuperJANET CONS Pilot Group

Taiwan's ATM Networks

- TNJCnet

Zeus Project

- Project Zeus

Back to Table of Contents

Appendix B - Annotated bibiliography.

Articles

1. Tokizawa, Ikuo; Kikuchi, Katsuaki,*"ATM link system technologies-super high-speed information highway for future multimedia communication services"* ,NTT Review vol 6 n 1 Jan 1994. p 49-55.

2. Cheng, Wood-Hi; Bechtel, James H.*"Gigabit fiber optic link using compact disc lasers"*, The International Soc. for Optical Eng. vol 2148 1884. p 210-215

3. Dumortier, Philip; Van Hauweermerien, Luc; Boerjan, Joost,*"Transport of gigabit ATM cell streams over lower order SDH backbone"*, Proc. - IEEE INFOCOM v 3 1994. IEEE, Piscatway, NJ p 1160-1169

4. Watson, Greg; Banks, David; Calamvokis, Costas; Dalton, Chris,*"AAL5 at a Gigbit for a kilobuck"*, Jrnl of High Speed Networks v 3 n 2 1994. p 127-145

5. Bohm, Christer; Lindgren, Per;*"DTM Gigabit Network"*, Jrnl of High Speed Networks v 3 n 2 1994. p 109-126

6. Morrison, John*"Gigabit Network Applications at Los Alamos National Lab."*, Conf. on Optical Fiber Communication, Tech Digest Series v 4 1994. p 64-65

7. Antonio, John K.*"Concurrent communication in high-speed wide area networks"*,IEEE Trans. on Paralled and Distributed Systems v 5 n 3. Mar 1994 p 264-273.

This is a theoretical paper. A parameter called receptivity is formally defined which is a measure of the amount of concurrent communication supported by the source-destination pairs in a high-speed network. The simulation results agree well with the theoretical predicitions. It also gives a easily computable approximations for *receptivity*.

8. Zurfluh, E.A; Cideciyan, R.D,*"IBM Zurich Research Laboratory's 1.13 Gb/s LAN/MAN prototype"*, Computer Networks and ISDN Systems v 26 n 2 Oct 1993. p 163-183

9. Joseph D. Touch,*"Defining High-Speed Protocols: Five Challenges and an Example That Survives the Challenges"*, USC/Information Sciences Institute, August 1994.

In this paper the author demonstrates how the WWW service meets the five challenges in high-speed networks and also meets the defining characteristics of a application requiring high-speed bandwidth.

10. Ahmed Tantawy, Odysseas Koufopavlou, Martina Zitterbart and Joseph Abler, *"On the Design of a MultiGigabit IP Router"*, Journal of High Speed Networks, vol 3 (1994), pg 209-232.

Routers will form the main bottleneck in high speed computing environment. This paper shows that fast IP routers can be built by identifying time critical points in the data flow and developing a special module for processing those points.

11. I. Chlamatc, A. Fumagalli, L. G. Kazovsky and P. T. Poggiolini,*"A Contention/Collision Free WDM Ring Network for Multi Gigabit Packet Switched Communication"*, Journal of High Speed Networks, vol 4 (1994), pg 201-219.

12. Abhaya Asthana, Catherine Delph, H. V. Jagadish and Paul Krzyzanowski, *"Towards A Gigabit IP Router"*, Journal of High Speed Networks, vol 1 (1992), pg 281-288.

13. K. Y. Eng, et al.,*"A Prototype Growable 2.5 GB/s ATM switch for BroadBand Applications"*, Journal of High Speed Networks, vol 1 (1992) pg 237-253.

14. James P. G. Sterbenz and Gurudatta M. Parulkar,*"Design of a Gigabit Host-Network Interface"*, Journal of High Speed Networks, vol 2 (1993), pg 27-62.

15. Jean-Paul Nussbaumer, et al.,*"Networking Requirements for Interactive Video on Demand"*, Gigabit Networks Workshop, 1994.

In this paper the author has proposed different cost measuring parameters for estimating the cost of providing interactive video on demand service. Various caching and tree networks are considered for providing this service.

16. Ahmed E. Kamal, Bandula W. Abeysundara,*"A Survey of MAC Protocols for High-Speed LANs"*, High Performance Networks, 1994.

The authors have presented an extensive survey on set of MAC protocols for high-speed LANs. The different LAN technology are classified according to Timing, Topology and Access Mode and studied in view of supporting gigabit speeds.

URLs

1. BAGNet home page, *http://chocolate.pa.dec.com/bagnet.html*

Bay Area Gigabit Testbet homepage. Give lot of details about the testbed and there is alo a map showing the testbed connectivity.

2. Gigabit Networking
http://www.yahoo.com/Computers_and_Internet/Communications_and_Networking/Gigabit_Networking/

Very useful starting point for other references on Gigabit Networking.

3. HIGH-Speed Networking
http://www.yahoo.com/Computers_and_Internet/Communications_and_Networking/High_Speed_Networking/

Another useful starting point for other points on Gigabit Networking.

4. High Performance Networks and Distributed Systems Archive, *http://hill.lut.ac.uk/DS-Archive/index.html*

Useful starting point for other points on Gigabit Networking, which leads to pointer of various gigabit testbeds.

Books

Craig Patridge, "Gigabit Networking", Addision-Wesley 1993.

This is the first book on gigabit networking. The author has been a frontier researcher in the area of gigabit networking. The books gives a comprehensive overview of the gigabit networking research. The book is very readable and not very mathematical.

Organisations

1. National Coordination Office for HPCC (High Performance Computing and Communications).
2. The Corporation for National Research Initiatives

Back to Table of Contents

Other Reports on Recent Advances in Networking
Back to Raj Jain's Home Page ***Last Update Aug 23, 1995***

Raj Jain is now at Washington University in Saint Louis, jain@cse.wustl.edu **http://www.cse.wustl.edu/~jain/**