Current Issues and Trends in Optical Networking

a report by
Raj Jain *and* **Naveen Bisht**

Chief Technology Officer and Co-founder, and Chief Executive Officer, President and Founder, Nayna Networks, Inc.

New Address: Raj Jain, Washington University in Saint Louis, jain@cse.wustl.edu, http://www.cse.wustl.edu/~jain

Introduction

The telecoms industry experienced a rapid downturn in 2001. The metro market was expected to achieve strong growth. Institutional money flowed into this sector and created dozens of start-up competitors in every segment, each claiming its own position and striving to provide cheaper, faster and smaller products. However, the demand turned out to be less than was anticipated. In addition, with the obvious price competition that results in such a situation, a seemingly promising situation now appears discouraging. The following will also have a major impact on the industry in 2002 and ahead.

- Capital expenditure budgets for carriers are still declining by about 24% in 2002.
- Carrier spending is transitioning from large network buildouts to incremental bandwidth expansion, based on an immediate revenue generation model and also by extracting more efficiency out of the deployed capacity.
- Oversupply of optical components caused by consolidation and bankruptcy of some of the carriers is expected to last through at least the first half of 2002.
- There have been spending disruptions due to carrier consolidation, management changes, reorganisations etc.
- Pricing pressures exist at all levels, from systems and components to services.

A silver lining, however, exists among this negativity. Optical systems based on new technologies will offer new services opportunistically at a reasonable cost by incorporating new functionality. Much of the innovation is required to enable carriers to offer new, higher-margin services and reduce their capital and operational expenditures and network costs. This will be accomplished by innovation at the component level, at the systems level by incorporating these innovative components with intelligent software, and at the architectural level by leveraging new architectures to result in cost savings and higher margins.

Ultimately, the end-user customer will drive growth in the industry. The days of getting the product to market fast are over, and the new perspective is success-based, revenue-based and incremental investment in the network. On the business side, lambda services are gaining traction with the carriers, and these services aim to provide line functionality at high speed and with competitive costs. Although the outlook for 2002 may still not be terribly optimistic for the optical networking industry, the foundation for innovation has been set. This innovation offers a compelling value proposition for the carriers to use as they start planning their next phase of network upgrade, buildout and maximising the efficiency of their networks.

In spite of the economic downturn, it is important to note that the number of Internet users has increased consistently in the past five to six years. In the late 1990s, Unix-to-Unix Network (UUNET) traffic was almost doubling every 90 days. While the current rate of growth may not be as high, it is still growing at a rapid pace. This growth will provide opportunity for the system vendors and component vendors to be able to leverage their innovative products to build into systems for the carrier infrastructure upgrade that they will have to implement in order to remain competitive in the marketplace.

Each technology follows an S-shaped curve as shown in *Figure 1a*. If the number of problems solved as a function of time required to solve them using a given technology were plotted, the curve would have an S shape. Initially, when a technology is just evolving, it takes some time to solve very few problems. In financial terms, it means that substantial funds are required by the technology developers but the resulting commercial value may be little. This phase of the technology is called the 'research phase' and is funded by governmental research funding organisations such as the National Science Foundation (NSF) or the Defense Advanced Research Project Agency (DARPA) in the US. After the fundamental problems have been solved, the technology growth curve takes an upturn and problems are solved quicker. The technology is then taken over by the commercial sector. Finally, after all 'easy' problems have been solved, the technology growth curve takes another turn. Now, only difficult problems are left and it takes considerable time to solve very few problems. In other words, the technology is no longer cost-effective. At this point, the researchers and the commercial world move on to newer technology.

Computer networking has also followed an S-shaped curve. If either the number of hosts on the Internet or the total number of Internet users were plotted, the curve would be S shaped.

The technology started with the concept of packet switching in 1968 and the initial implementation of the four-node Advanced Research Projects Agency Network (ARPAnet) in 1969. The 'knee' of the networking growth curve occurred in 1995 with the popularity of the World Wide Web. The peak growth rate (the centre point of the S curve) happened sometime in late 1999 or early 2000. At this point, we are slightly beyond the centre point as shown in *Figure 1b*. This is true for the electronic networks. Conversely, optical networking is still near the knee. There is quite a bit of research activity here and the commercialisation of the optical networking technologies is just starting to take place. An exponential growth in optical networking should be witnessed soon.

There are four issues that are being hotly debated this year at networking conferences. These are:

1. bandwidth glut versus traffic growth;
2. all-optical (OOO) versus optical-electrical-optical (OEO);
3. Ethernet versus Synchronous Optical Network (SONET); and
4. mesh versus ring.

Bandwidth Glut versus Traffic Growth

One of the fundamental issues that is being debated is whether network traffic growth has stopped. The starting point of this debate was a forecast by McKinsey & Co and JPMorgan in May 2001 that Internet traffic growth will slow down from 200% to 300% per year to 60% by 2005. Soon after that, the *Wall Street Journal*, the *New York Times*, *Forbes* and other popular media papers and magazines reported the finding that 98% of fibre is unlit. Nortel blamed the loss of its revenue on the falling Internet Protocol (IP) traffic. A Merrill Lynch analyst reported that carriers are using only an average of 2.7% of their total lit fibre capacity.

Figure 1: Technology Growth Curves

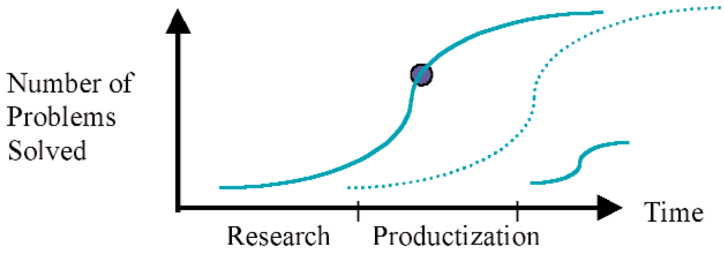

(a) General Technology Growth Curve

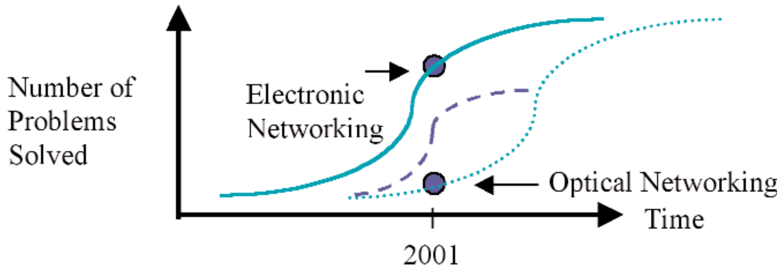

(b) Electronic vs Optical Networking Growth Curve

All these discouraging statements and forecasts accentuated the fall of the telecoms market. However, when analysing these statements, although they might be true, there should not be cause for alarm. Of the three statements mentioned previously, the first is a forecast while the other two are based on current usage. While the forecast may or may not be true, the current usage needs to be explored further. Installed fibre alone does not constitute the telecoms infrastructure. Fibre installation costs are high compared with the cost of the fibre itself and so it is quite common for carriers to install cables with hundreds of fibres when they need just a few. Having unlit fibre is also quite common and should not be considered alarming. Similarly, average usage of networking facilities is generally very low. For example, most computers today are equipped with 100 megabit per second (Mb/s) Ethernet ports. The average usage of these ports is very low – generally, less than 1%. Nevertheless, networking links are often bottlenecks and we find ourselves waiting for the information to arrive over the network. At these peak usage times, we need all 100Mb/s and more. Thus, the network capacity is planned for peak usage and not the average.

Telechoice – a market research company – conducted a usage study for Williams Communications. They studied the 22 most used routes in the US and found that the utilisation on 14 of these routes exceeded 70%. This high utilisation would necessitate upgrading these routes and the networking equipment. Larry Roberts, one of the early founders of the Internet, himself conducted a measurement study of traffic on 19 of the largest Internet service providers (ISPs) in the US. He

Figure 2: Measured Traffic Growth of 19 Largest-tier US ISPs

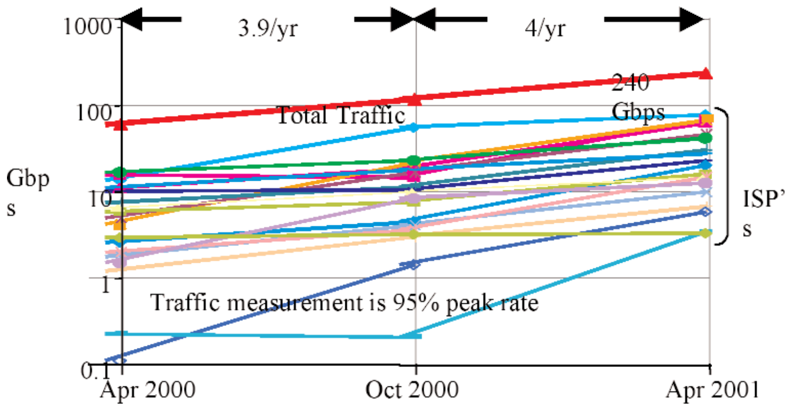

Source: Lawrence G Roberts

Figure 3: Traffic Growth versus Electronic Circuit Speed Growth

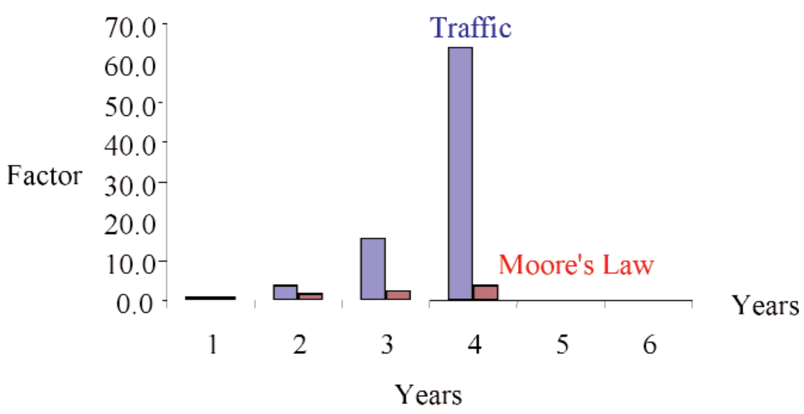

obtained data on 95% usage for each of these ISPs at six-month intervals between April 2000 and April 2001. A plot of this data is shown in *Figure 2*. The figure also shows 95% of the total traffic on the measurement dates. Roberts concluded that the traffic growth factor was 390% per year during April 2000 to October 2000 and it was 400% during the following six months. This study has been very helpful in calming down the bandwidth glut debate, at least for the time being.

OOO versus OEO

Gordon Moore, one of the founders of Intel, observed that the number of transistors in integrated circuits doubles every 18 months. This observation has been found to be true over the last two decades and is often interpreted to mean that the speed of electronic circuits is doubling every 18 months. Even this high rate of growth is slow when compared with the growth of the network traffic. *Figure 3* shows the growth of network traffic and the speed of electronic circuits as a function of time. The traffic growth has been assumed to be 400% per year using the Roberts study mentioned previously. Note that within a few years, the network traffic growth is expected to be several orders of magnitude greater than what can be handled by electronic circuits without parallelism. This is an argument in favour of optical switching.

In current switches, even though the signal comes in optically on a fibre and goes out optically on a fibre, it must be converted to electronic form for switching. This is known as OEO switching. In OOO switches, the signal is switched optically. There are several methods for optical switching, the most common being the use of micromirrors with microelectromechanical systems (MEMS). Based on the switching instructions, the mirrors can be rearranged quickly so that the light coming in on one fibre is diverted to another fibre as desired.

OOO switches are data format independent in the sense that the data being switched could have any data link format such as SONET, Ethernet, synchronous digital hierarchy (SDH) or optical transport network (OTN) (based on International Telecommunication Union (ITU) G.709). The same OOO hardware can support all data link formats. An OEO switch, on the other hand, will require a different circuit or software to support each of these formats.

The OOO switches are also relatively rate independent. The same hardware can switch a 2.5 gigabit per second (Gb/s) or 10Gb/s signal. Of course, the noise tolerances are tighter for 10Gb/s than 2.5Gb/s and so the component quality has to be better for 10Gb/s equipment than that designed for 2.5Gb/s. OEO switches, on the other hand, are highly rate dependent. An integrated circuit designed for 2.5Gb/s cannot handle 10Gb/s. In the simplest approach, four 2.5Gb/s circuits will be required to handle a 10Gb/s signal and so the cost of OEO switches grows proportionally to the data rate. The same argument applies for space and power. This is the prime reason in favour of OOO switching at high data rates. The conventional wisdom at this time is that, at 10Gb/s and higher rates, per-port cost of OOO switches is less than OEO switches.

The rate independence of OOO switches also reflects in their upgradability. Upgrading a 2.5Gb/s OOO switch to 10Gb/s may require fewer changes than an OEO switch. As mentioned previously, for an OOO switch, the same basic design can be upgraded with higher-quality components to support the higher rates. Upgrading an OEO switch, on the other hand, will require a complete change of design.

The following two issues are related to the ability to handle part of a wavelength or multiple wavelengths. An OEO switch can easily separate out a part of a signal, and so adding or dropping a part of the wavelength is easily accomplished. In an OOO switch, this can be accomplished only if the different parts of the signal have some differences in optical characteristics, such as time (slots), frequency, phase or polarisation. On the other hand, an OOO switch can switch a multiwavelength signal as easily as a single wavelength, provided the optical components are designed properly. OEO switches will require separate circuits to handle each of the wavelengths and would be very costly for the large number of wavelengths that can be accommodated in a fibre.

In terms of performance monitoring, OOO switches are bit-rate and format independent and so cannot see bit errors (rate or format errors) easily. OOO switches monitor optical defects such as wavelength shifts, optical signal-to-noise ratios or power levels. These defects also result in bit errors but not all bit errors are visible to optical monitors. To monitor electrical signal, OOO switches provide optional electrical monitoring. In OEO switches, the cost of monitoring is already built in since they have to verify the data rate and format before they can switch the signal.

If two or more signals need to be sent on a single fibre, they should have different optical characteristics (wavelength or polarisation). Therefore, it may sometimes be necessary to convert these optical characteristics (primarily the wavelength) at switching points. Optical approaches for wavelength conversion are a few years away. Several companies have announced products in this area but none are shipping at this point. Therefore, OOO switches offer optional OEO-based wavelength conversion on selected channels as needed. Often, this is not necessary since the dense wavelength division multiplexing (DWDM) equipment, which is used with both OEO and OOO switches, already has transponders that change wavelengths as required.

A comparison of OEO and OOO switches is summarised in *Table 1*. In summary, OOO switches are the direction for the future as we transition into higher speeds and larger switch capacities.

Table 1: OEO versus OOO

Feature	OEO	OOO
Data format dependence	*Yes*	*No*
Cost/space/power independent of rate	*No*	*Yes*
Upgradability to higher rate	*No*	*Yes*
Subwavelength switching	*Yes*	*Future*
Waveband switching	*No*	*Yes*
Performance monitoring	*Bit error rate*	*Optical signal degradation*
Wavelength Conversion	*Built in*	*Currently electronic*

Ethernet versus SONET Debate

It is now well established that the traffic on most carrier networks is predominantly data traffic. SONET/SDH technology is designed for voice traffic and is very expensive compared with Ethernet, which is designed for data. It is clear that using Ethernet switches in place of SONET add/drop multiplexers (ADMs) will reduce the cost considerably. However, there are several obstacles to the adoption of Ethernet technology, the primary obstacle being its reliability and availability.

SONET technology was designed primarily for carrier networks and has very robust reliability and availability mechanisms built in. In particular, SONET networks are designed to provide 99.999% availability, which is equivalent to a downtime of five minutes per year. This is achieved by a high level of redundancy inside and outside the equipment. Ethernet technology, on the other hand, was designed primarily for enterprise networks where availability requirements are not as high.

Recent interest by carriers in Ethernet is visible from the activities in the 10Gb/s standardisation efforts. 10G Ethernet is designed for two data rates: 10Gb/s for local area networking (LAN) applications and 9.5Gb/s for wide area networking (WAN or telecoms) applications. The WAN version uses SONET framing. It is compatible with SONET equipment except for the clock jitter requirements. SONET requires a clock jitter of 4.6 to 20 parts per million (ppm), while 10G Ethernet requires only 100 ppm. This decision was highly debated by the Institute of Electrical and Electronics Engineers, Inc. (IEEE) standards committee and may have resulted in the delay of the standard, but requiring tighter tolerances would have increased the cost of the equipment significantly. As a result, a 10G Ethernet signal cannot be sent directly to a legacy SONET ADM. Ethernet line termination equipment is required to buffer the incoming signal and send out the well-conditioned signal to the SONET equipment. In this way, the extra cost of clock jitter conditioning is not incurred if the entire WAN network is based on 10G Ethernet technology. This is the future plan of many carriers, particularly those that are primarily data carriers.

SONET networks are traditionally organised in dual-ring topologies that allow for a very fast recovery from node and link failures. Ethernet equipment is traditionally organised as mesh networks. Ethernet switches use a spanning tree algorithm to automatically convert the mesh topology to a spanning tree topology for forwarding. The spanning tree takes a few minutes to converge and so restoration times may be long. IEEE has designed a rapid spanning tree algorithm.

In order to provide fast recovery time for Ethernet traffic, the IEEE has also commenced a Resilient

Table 2: SONET versus Ethernet

Feature	SONET	Ethernet
Bit-rate (b/s)	155M, 622M, 2.5G, 10G, 40G...	1M, 10M, 100M, 1G, 10G...
Timing	Isochronous (Periodic 125μs)	Plesio-isochronous
Multiplexing	Bit	Packet
Clocks	Common	Independent
Clock jitter	4.6–20 ppm	100 ppm (may change)
Usage	Telecoms	Enterprise
Volume	Millions	100s of Millions
Price (10Gb/s)	>10k	<1k
Recovery	50ms	Few minutes
Topology	Rings	Mesh

Packet Ring (RPR) project to allow Ethernet traffic to be sent on dual-ring networks. This will provide fast recovery times matching those of SONET, while, at the same time, being more efficient in terms of redundancy by allowing both rings to be used when there is no failure. A comparison of SONET versus Ethernet is summarised in *Table 2*.

Now that the 10G Ethernet standards are nearing completion, the IEEE has started discussing the next steps. In particular, a survey of IEEE 802.3ae members has indicated that 70% of the members would like the next version of the Ethernet to be 40Gb/s rather than 100Gb/s. This will allow the OC-768 technology being developed for SONET to be reused for Ethernet.

While Ethernet is being modified to take on the best features of SONET, SONET is also undergoing changes to become better suited for data traffic. The main problems of legacy SONET are as follows.

- Line rates are highly discrete. The only available rates are synchronous transport signal (STS)-3c, STS-12c and STS-48c. For intermediate rates, the customers are required to get the next higher STS rate, which may be too high.

- Rates of STS-3C (155Mb/s), STS-12c (622Mb/s), STS-48c (2.4Gb/s) do not match well with those of the data traffic that originates mostly from Ethernet LANs. Thus, to connect two 100Mb/s Ethernet LANs, users are forced to use STS-3c and waste one-third of the bandwidth.

- Multipath and traffic splitting are not supported. To send 100Mb/s traffic, for example, two 50Mb/s (or STS-1) paths cannot be used.

- SONET streams are fixed rates while the data is highly bursty. It is not possible to change dynamically the rate allocated to a customer.

- STS-48c rate is very close to the physical layer rates of two gigabit Ethernet LANs but falls short by a mere 100Mb/s. Gigabit Ethernet uses 8b/10b encoding at the physical layer and so to interconnect two remote Ethernet LANs at the physical layer using SONET will require a payload rate of 2.5Gb/s. This slight mismatch forces carriers to allocate one full STS-40c channel for each gigabit Ethernet stream.

- Payload type is indicated in one of the overhead bytes and so each SONET frame can only have one type of payload. For example, if the C2 (path signal label) byte in the SONET path overhead is 22, the entire payload should consist of scrambled Point-to-Point Protocol (PPP) frames. It is not easy to mix different types of frames in one SONET frame.

The next generation of SONET has solved all of these problems as follows.

- Virtual concatenation allows any number of STS-1s or T-1s to be grouped as one SONET stream. For example, two STS-1s can be combined to form a STS-1-2v stream that has a rate of 102Mb/s suitable for sending 100Mb/s Ethernet frames.

- Virtual concatenation also allows component streams to take different paths. For example, it will allow a carrier to offer OC-192 services using four parallel OC-48 paths.

- With STS-1 and T-1-level virtual concatenation, it is possible to match very closely any required data rate.

- A link capacity adjustment scheme (LCAS) allows the number of STS-1s in the virtual concatenated SONET stream to be changed dynamically.

- Generic Framing Protocol (GFP) allows each packet in the SONET frame to have its own protocol type and so it is possible to transport frame relay, fibre channel and Ethernet in the same SONET stream.

- GFP also has a transparent mode that is designed to compress 8b/10b encoded streams by a factor of 80/65. Thus, one gigabit Ethernet stream requires only 1.02Gb/s for physical layer connectivity. This allows two gigabit Ethernet streams to be carried easily in one STS-48c link.

In summary, the next-generation SONET with GFP, LCAS and virtual concatenation will be better suited for data streams and may provide a good compromise in terms of flexibility and availability. There is no question that Ethernet will continue to

Figure 4: Ring versus Mesh Networks

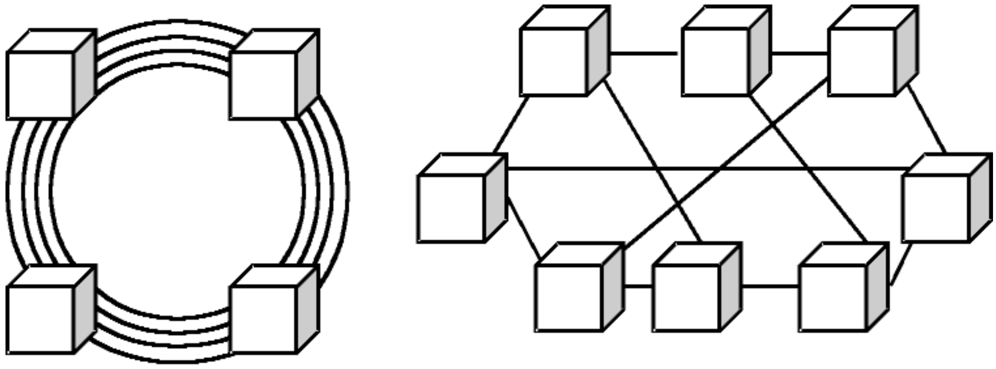

be the most popular data service, but the question still remains whether Ethernet should also be used as a transport or whether the Ethernet frames can be transported using a SONET infrastructure. In the current telecoms world dominated by incumbent local exchange carriers, the possibility of a SONET infrastructure with next-generation features is more promising than that of Ethernet transport.

Ring versus Mesh

Telecoms networks are currently organised in ring-based topologies while the data networks use mesh-based topologies. This has started the debate on the merits of ring and mesh topologies. This debate is similar to that between Ethernet and token rings back in the initial days of IEEE 802 standardisation.

As shown in *Figure 4*, on a ring, all links have to have the same data rate. If any one link is upgraded, for example from 2.5Gb/s to 10Gb/s, all links and nodes have to be upgraded to that rate. Therefore, rings are more suitable for networks where the traffic between the nodes is homogeneous. Rings are not generally used in long-haul networks where the traffic is highly non-homogeneous. The mesh networks, on the other hand, allow incremental upgrades. Any link can be upgraded to a higher rate while the others can remain at a lower rate. Similar arguments apply in the case of DWDM networks where the number of wavelengths must be the same on all nodes of a ring.

Mesh networks typically require 50% less protection and 50% less working capacity than rings. This is because of the inherent spatial reuse feature of mesh networks, whereby the links not being used by one flow can be used by other flows. The savings from mesh networks increase as the degree of connectivity increases and as the non-homogeneity of the traffic increases.

Currently, there are two parallel efforts in the telecoms networking community. On one side, they are trying to develop mesh-based protection and restoration mechanisms and protocols, while, on the other side, they are also trying to develop ring-based mechanisms and protocols suitable for data traffic. This second effort is reflected in the RPR work by the IEEE. RPR uses a dual-counter rotating ring topology similar to that used in SONET and Fibre Distributed Data Interface networks.

Summary

The optical networking industry has been affected by the economic slowdown. However, we believe that the slowdown is temporary. Internet traffic is still growing. Carriers need to find ways to increase their revenues from this increasing traffic and reduce capital and operating expenditures. Optical cross-connects and mesh topologies seem to be more economical than their counterparts. Ethernet services using a next-generation SONET infrastructure are now more likely than an entirely Ethernet-based infrastructure.