

Optical Networking: Recent Advances, Trends, and Issues

Naveen Bisht, CEO, President and Founder
Raj Jain, CTO and Co-Founder
Nayna Networks, Inc.
461 Sycamore Drive
Milpitas, CA 95035
Email: Naveen@nayna.com, Raj@nayna.com

Abstract

This paper summarizes the latest optical telecom networking market developments. Issues that are being debated hotly in the networking community are presented. This includes the issue of traffic glut vs growth, optical vs electrical switching, Ethernet vs SONET, and mesh vs ring. We then also explain the trend towards optical crossconnects.

Introduction

The telecommunications industry experienced a rapid downturn in 2001. As a number of industry gurus attempted to “call the bottom,” some of the most generally accepted beliefs turned into mirages. Take a look at the metro market. It was believed that the metro market was poised for strong growth. Institutional money flowed into this sector and created dozens of startup 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 smaller than it was anticipated. On top of that, with the obvious price competition that results in such a situation, a seemingly great opportunity now looks significantly gloomier. The question, which comes to the mind of every telco professional, is “Have we hit the bottom yet?” Perhaps, not. There are a few factors, mentioned below, which will have a large impact in 2002 and ahead.

1. Capital expenditures budgets for carriers are still declining about 24% in 2002.
2. Carrier spending is transitioning from large network build outs to incremental bandwidth expansion, based on immediate revenue generation model and also, by extracting more efficiency out of the deployed capacity.
3. Oversupply of optical components caused by consolidation, bankruptcies of some of the carriers are expected to last through at least first half of 2002.
4. Spending disruptions due to carrier consolidation, management changes, reorganizations etc.
5. Pricing pressures at all levels of the food chain from systems, components to services.


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

Ultimately the end user customer, whether it's business or residential, will drive growth in our industry. The days of “Get to market fast, build it and customers will come with open arms” are gone and the mantra going forward 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 competitive costs. Although 2002 may still not look very good for us in the industry, the foundation for innovation in the optical networking industry 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, build out and maximizing the efficiency of their networks.


In spite of the economic downturn, it is important to note that the number of Internet users have grown consistently in past 5-6 years. In the late 90's, UUNET traffic was almost doubling every 90 days. While the current rate of growth may not be as high as during those times, 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 which they will have to do in order to remain competitive in the market place.

Each technology follows an S-shaped curve as shown in Figure 1a. If we were to plot the number of problems solved as a function of time required to solve them using a given technology, the curve would have an S shape. In the beginning when a technology is just evolving, it takes quite a bit of 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 organizations such as the National Science Foundation (NSF) or the Defense Advanced Research Project Agency (DARPA) in the United States. After the fundamental problems have been solved, the technology growth curve takes an upturn and it takes very little time to solve many problems. In other words, it takes very little money to produce quite a bit of money using that technology. 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 hard problems are left and it takes quite a bit of 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 some newer technology.



(a) General Technology Growth Curve



(b) Electronic vs Optical Networking Growth Curve

Figure 1: Technology Growth Curves

Computer networking has also followed an S-shaped curve. If we were to plot either the number of hosts on the Internet or the total number of Internet users, we would find that the curve is S shaped. The technology started with the concept of packet switching in 1968 and the initial implementation of 4-node ARPAnet in 1969. The knee of the networking growth curve occurred in 1995 with the popularity of World Wide Web. The peak growth rate (the center point of the S curve) happened sometime in late 1999

or early 2000. At this point, we are slightly beyond the center 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 going on and the commercialization of the optical networking technologies is just starting to take place. We should soon see an exponential growth in optical networking.

There are 4 issues that are being hotly debated this year at networking conferences. These are:

1. Bandwidth Glut vs Traffic Growth
2. OOO vs OEO
3. Ethernet vs SONET
4. Mesh vs Ring

We explain both sides of these debates next in this paper and also explain the trend towards optical crossconnects.

1. Bandwidth Glut vs Traffic Growth

One of the fundamental issues that are being debated is whether network traffic growth has stopped. The starting point of this debate was a forecast by McKinsey & Co and J.P. Morgan in May 2001 that Internet traffic growth will slow down from 200-300% per year to 60% by the year 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 the fiber is unlit. Nortel blamed the loss of its revenue on the falling IP traffic. Michael Ching, a Merrill Lynch analyst, reported that carriers are using only average 2.7% of their total lit fiber capacity.

All these gloomy statements and forecasts accentuated the fall of the telecom market. However, when one analyzes these statements, one finds that while these statements might be true, they should not be a cause for alarm. Of the three statements mentioned above, the first one 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 fiber alone does not constitute the telecom infrastructure. Actually, fiber installation cost are high compared to the cost of the fiber itself and so it is quite common for carriers to install cables with hundreds of fibers when they need just a few fibers. Having unlit fiber is quite common and should not be considered alarming. Similarly, average usage of networking facilities is generally very low. For example, most of the computers today are equipped with 100 Mbps Ethernet ports. The average usage of these ports is very low, generally, less than 1%. Nevertheless, networking links are often the bottlenecks and we find ourselves waiting for the information to arrive over the network. At those peak usage times, we need all 100 Mbps 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 United States and found that the utilization on 14 of these routes exceeded 70%. This high utilization 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 United States. He obtained data on 95-percentile usage at each of these ISPs at six months intervals between April 2000 and April 2001. A plot of this data is shown in Figure 2. The figure also shows 95-percentile of the total traffic on the measurement dates. Roberts concluded that the traffic growth factor was 390% per year during the April 2000 to October 2000 and it was 400% during the next six months. This study has been very helpful in calming down the bandwidth glut debate at least for the time being.

2. OOO vs 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 to 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 above. 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 favor of optical switching.




Figure 2: Measured Traffic Growth of 19 Largest Tier 1 US ISPs
(Source: Lawrence G. Roberts)

In current switches, even though the signal comes in optically on a fiber and goes out optically on a fiber, it must be converted to electronic form for switching. This is called Optical-Electrical-Optical or OEO switching. In Optical-Optical-Optical (OOO) switches, the signal is switched optically. There are several methods for optical switching, the most common being the use of micro mirrors using micro-electrical machines (MEMs). Based on the switching instructions, the mirrors can be rearranged quickly to make the light coming on one fiber go to another fiber 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 or Ethernet or SDH or OTN (based on 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 2.5 Gbps signal or 10 Gbps signal. Of course, the noise tolerances are tighter for 10 Gbps than 2.5 Gbps and so the component quality has to be better for 10 Gbps equipment than those designed for 2.5 Gbps. OEO switches on the other hand are highly rate dependent. An integrated circuit designed for 2.5 Gbps cannot handle 10 Gbps. In the simplest approach, four 2.5 Gbps circuits will be required to handle 10 Gbps 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 favor of OOO switching at high data rates. The conventional wisdom at this time is that at 10 Gbps and higher rates, per port cost of OOO switches is less than that of OEO switches.

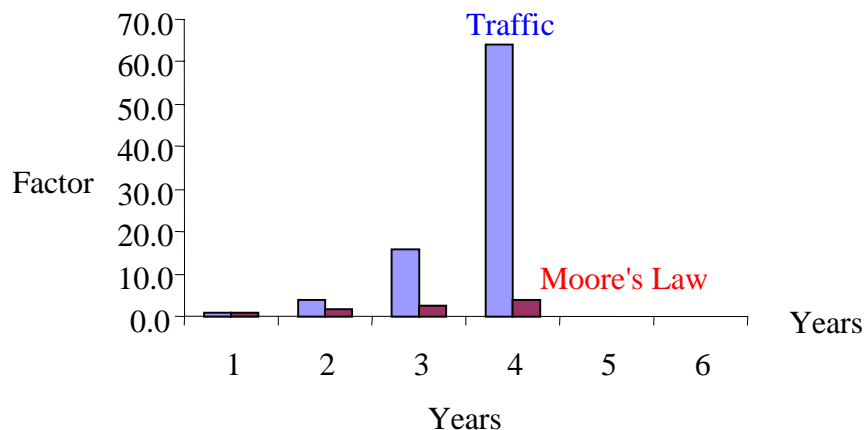


Figure 3: Traffic Growth vs Electronic Circuit Speed Growth

The rate independence of OOO switches also reflects in their upgradability. Upgrading a 2.5 Gbps OOO switch to 10 Gbps may require fewer changes than those in an OEO switch. As mentioned earlier, 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 next 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 done only if the different parts of the signal are different in some optical characteristics, such as time (slots), frequency, phase, or polarization. On the other hand, an OOO switch can switch a multi-wavelength signal as easily as a single wavelength provided the optical components are properly designed. 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 fiber.

In terms of performance monitoring, OOO switches are bit rate and format independent and so cannot easily see bit errors (which are basically rate or format errors). 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 fiber, they should have different optical characteristics (wavelength or polarization). 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 space 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 needed since the 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 summarized in Table 1.

Table 1: OEO vs OOO

Feature	OEO	OOO
Data Format Dependence	Yes	No
Cost/Space/Power independent of rate	No	Yes
Upgradability to higher rate	No	Yes
Sub-Wavelength Switching	Yes	Future
Waveband Switching	No	Yes
Performance Monitoring	Bit error rate	Optical signal degradation
Wavelength Conversion	Built-in	Currently Electronic

In summary, OOO switches are the direction for the future as we transition into higher speeds and larger switch capacities.

3. Ethernet vs 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 to Ethernet, which is designed for data. It is clear that using Ethernet switches in place of SONET ADMs will reduce the cost considerably. However, there are several obstacles to the adoption of the Ethernet technology, the primary one being the 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% (five nine's) availability, which is equivalent to a down time of 5 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 that high.

Recent interest of carriers in Ethernet is visible from the activities in the 10 Gbps standardization efforts. 10 G Ethernet is designed for two data rates: 10 Gbps for local area networking (LAN) applications and 9.5 Gbps for wide area networking (WAN or telecom) 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 10 G Ethernet requires only 100 ppm. This decision was highly debated in the 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 of this, 10 G Ethernet signal cannot be sent directly to a legacy SONET ADM. An "Ethernet Line Termination Equipment (ELTE)" is required to buffer the incoming signal and send out the well-conditioned signal to the SONET equipment. This way the extra cost of clock jitter conditioning is not incurred if the entire WAN network is based on 10 G Ethernet technology. This the future plan of many carriers, particularly those that are primarily data carriers.

SONET networks are traditionally organized in dual ring topologies that allow for a very fast recovery from node and link failures. Ethernet equipment is traditionally organized 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 the restoration times may be large. IEEE has started an effort to design a fast spanning tree algorithm.

In order to provide fast recovery time for Ethernet traffic, IEEE has also started a "Resilient Packet Ring" (RPR) project which will 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 vs Ethernet is summarized in Table 2.

Table 2: SONET vs Ethernet

Feature	SONET	Ethernet
Bit Rate (bps)	155 M, 622 M, 2.5 G, 10 G, 40 G, ...	1M, 10 M, 100 M, 1 G, 10 G, ...
Timing	Isochronous (Periodic 125 μ s)	Plesio-Isochronous
Multiplexing	Bit	Packet
Clocks	Common	Independent
Clock jitter	4.6 to 20 ppm	100 ppm (May change)
Usage	Telecom	Enterprise
Volume	Millions	100's of Millions
Price (10 Gbps)	>10k	<1k
Recovery	50 ms	Few Minutes
Topology	Rings	Mesh

Now that the 10 G Ethernet standards are nearing completion, 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 40 Gbps rather than 100 Gbps. This will allow the OC-768 technology being developed for SONET to be reused for Ethernet.

4. Ring Vs Mesh

Telecommunications networks are currently organized in ring-based topologies while the data networks use the 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 standardization.

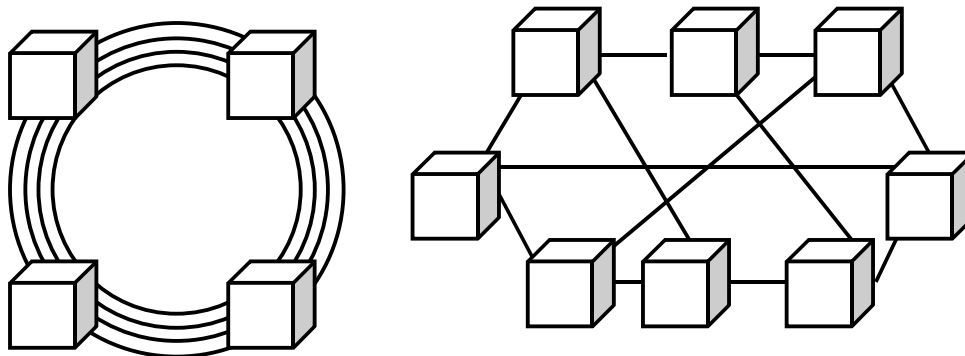


Figure 4: Ring vs Mesh Networks

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 2.5 Gbps to 10 Gbps, all links and nodes have to be upgraded to that rate. Therefore, rings are more suitable for networks where the traffic among 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 telecom 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 "Resilient Packet Rings" (RPR) work going on in IEEE. RPR uses a dual counter rotating ring topology similar to that used in SONET and FDDI networks.

5. Optical Crossconnects

As the network traffic increases, the capacity of crossconnects and switches must also increase. There are two ways to increase the capacity: higher port count or higher speed per port. For example, 80 Gbps traffic can be handled with eight 10G ports or two 40G ports. Increasing the port speed increases the number of multiplexing and de-multiplexing (grooming) points and increases the number of hops. This has led to a trend toward larger number of ports per switch. A similar argument applies towards the number of wavelengths and the data rate per wavelength as well.

A single crossconnect with a large number of ports is more cost effective than an interconnection of several smaller crossconnects. To illustrate this, consider the two alternatives shown in Figure 5 for a 24×24 crossconnect. Alternative A consists of using 24 4×4 switches with a total of 192 ports. The alternative B consists of a single 24×24 switch with a total port count of 48. Since most of the cost of the cross connect is in the port cost, alternative A is almost 4 times more expensive than alternative B. In general, larger crossconnects are more cost effective than multiple smaller ones. This has led to a trend towards large crossconnects, which can be more economically realized with optical technology than electrical.

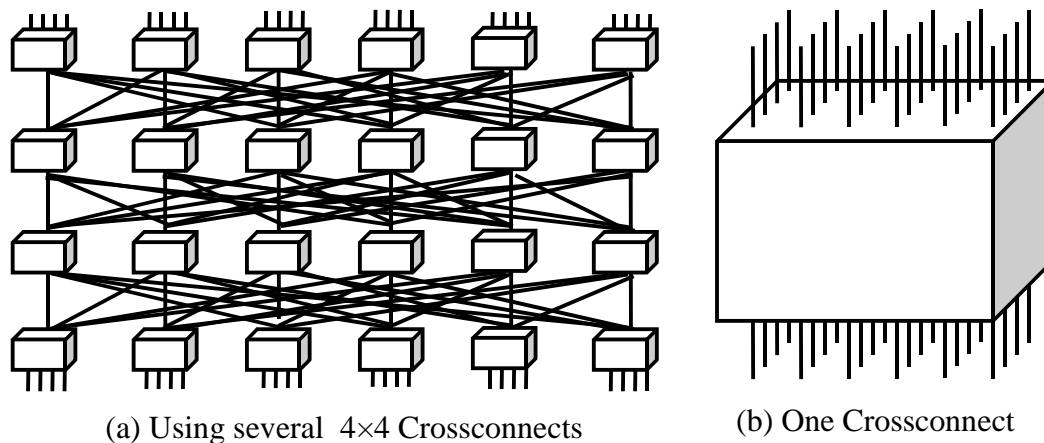


Figure 5: Two ways to synthesize a 24×24 Crossconnect

Summary

Optical networking industry has been affected by the economic slowdown. However, we believe that the slowdown is temporary. The Internet traffic is still growing. Carriers need to find ways to increase their revenues from this increasing traffic and reduce their capital and operating expenditures. Optical crossconnects, larger number of ports, mesh topologies, and Ethernet based networks seem to be more economical than their counterparts.