

Report of the Breakout Session on “Architectures and Protocols for Optical Networks” Held During the NSF ANIR Principal Investigator (PI) Meeting in January 2003.

Summarized by **Chunming Qiao & Byrav Ramamurthy**
with contributions from the participants listed below

Participants of the Breakout Session

Chunming Qiao (Lead), Byrav Ramamurthy (Scribe),
Ken Christensen, S.L. Chuang, Rene Cruz, Arjan Durresi, Y. Fainman, Andrea Fumagalli, Larry
Gallery, Wendy Huntoon, Julio Ibarra, Jason Jue, Admela Jukan, Prem Kumar, Jason Leigh, Joe
Mambretti, Eytan Modiano, Biswanath Mukherjee, Peter O’Neil, Martin Reisslein, Arun
Somani, Malathi Veeraraghavan, Paul K. Yu.

Introduction

It is generally understood that an optical network can be either all-optical (where both control and data remain in the optical domain at an intermediate node), photonic or almost all-optical (where only data remains in the optical domain at an intermediate node but control is done electronically), or even hybrid where both optical and electronic switching may take place at an intermediate node. Both electronic and photonic devices can use analog or digital switching. In addition, data may be switched in an optical network based on circuit switching, packet/burst switching, or some combination of the two using variable-length packets.

The Case for Optical Networking

The current downturn of the telecommunications industry in general, and optical networking sector in particular, is a result of its own success in providing high-capacity point-to-point bandwidth, which has increased by a factor of 1000 in the past decade. Yet despite an obvious demand this technical achievement has not yielded a corresponding increase in the data bandwidth available to businesses or private individuals. Improvements at the device level – such as a higher speed transmitter or detector – are not easily translated into improvements in a complex network. The practical difficulties in designing, administering, and maintaining a dynamic optical network remain, and are unlikely to be solved anytime soon. Still, optical networking, employing either all-optical or photonic switching, is the only core networking technology that can support dynamic bandwidth allocation and transport for the increasing amount of traffic expected in the foreseeable future.

This is partly due to the fact that the Internet traffic is still doubling almost every year. In addition, it is expected that a large number of users in scientific computing and academic communities, as well as in military and other government agencies will require Gbps or higher bandwidth for new and emerging applications such as Grid Computing and Tele-immersion that will likely also require minimum delay and jitter, and maximum security.

Meanwhile, despite the impressive progress made in electronics, electronic switching (or routing) is fast approaching its peak throughput of several terabits per second. The limitations of electronic switching are not fundamental but rather practical ones imposed largely by realistic bounds on power and space consumption. Although in general, electronic switching still enjoys a much better cost-performance ratio, and will likely remain more suitable for a majority of traffic in the near term, the long-term potential of optically switching a majority of traffic in the core at a much higher throughput with much less power and space consumption (and accordingly a lower cost) through innovative device, component and networking approaches cannot be ignored.

WDM transmissions carry up to 128 independent data signals on the same fiber. These optical signals must be demultiplexed into individual components for electronic switching. Providing receivers and transmitters for all the wavelength channels is a prohibitively expensive approach to gain access to a single wavelength signal, yet passive filters lack the dynamic capacity for bandwidth allocation. This has been one of the major barriers to dynamic networking.

Optical networks can sustain a much higher throughput than pure electronic switching/routing approaches. In particular, all-optical or photonic switching can provide virtually unlimited bandwidth. By providing all-optical bypass and add/drop, as well as all-optical/photonic switching, they also make the Internet more scalable compared to an approach using just point-to-point fiber (WDM) links directly between IP routers. In addition, optical networking makes the Internet more flexible because they can provision bandwidth (fat-pipes) between IP routers fast, thus enabling new services. Also, optical networks can not only support general-purpose networks such as the Internet, but also special-purpose networks e.g., Virtual Private Networks (VPNs) by establishing lightpaths (wavelength-routed connections) using a subset of wavelengths available on each fiber link. Finally, optical networks can provide a high survivability/availability through protection against failures such as fiber-cuts and fast recovery from failures and restoration of services. This is especially important as the transmission speed of a fiber-optic link continues to increase to Terabits per second and beyond, while at the same time, our nation's economy depends more and more on the communication infrastructure. Without optical protection and restoration, it would be extremely inefficient for IP to re-route the massive amount of data (up to hundreds of thousands of IP flows) affected by a single fiber cut.

Today, optical networks carry at least 90% of the total traffic (data and voice) in the core. Optical networks are also moving into the Metro area, and are a promising alternative for providing broadband access.

Why is Continued NSF Funding in Optical Networking Important?

Research problems in optical networking are “hard” because they truly require inter-disciplinary effort in enabling optical/photonic and electronic device and component technologies as well as networking level system architecture, hardware and software development. They also require integrated, cross-layer solutions taking into consideration new, emerging applications, TCP/IP and other protocols, SONET functionalities and WDM specific features. In addition, this area of research is still relatively young and fast growing. In general, there is a “chicken and egg” situation wherein advanced architectures might evolve based on novel component technologies,

and new components might be created to support futuristic architectures – this is why both networking research and component research need continuous support.

In the last decade or so, several US government funding agencies, notably DARPA, have spent a large amount of money on optical networking research. While good return on investment has been achieved, the fact is that the portion of the funds spent on networking level research is relatively small compared to that spent on transmission, and device/component technologies. NSF has also funded a few multi-disciplinary research projects but again, substantially more money has been spent on device/component level research than on networking research. The optical networking industry in general, and many start-up companies in particular, have also raised a large amount of money, most of which was also spent on device/component R&D than networking level research.

While device/component research has been and will remain crucial for the advancement of optical networking technologies and should be supported, optical networking research is useful for guiding the evolution of device/component technologies and is an inherently multi-disciplinary research field.. Even with today's device/component technologies, there are challenging networking research problems that need to be addressed. Among these are how to operate in the absence of wavelength converters (which are currently costly and technologically immature) without losing performance by using efficient routing and wavelength assignment (RWA) algorithms in circuit-switched networks and to develop innovative contention resolution methods in packet/burst switched networks. Another example of a networking level research problem, which is unique to optical technology and also fundamental in that it will not be solely decided by the device/component technology advances, is how to gracefully accommodate a mixture of optical circuit-switching (such as wavelength-routing), optical packet switching (OPS), and optical burst switching (OBS) techniques, to yield a well-managed and reliable network . Such a problem should be addressed as early as possible because, for instance, if OPS were found to be fundamentally less appealing than OBS, OCS or a combination of OBS and OCS, research into optical RAM would not make sense (but given that we are not certain of that, research into OPS and optical RAM should be supported).

It should also be noted that networking level research and device/component level research should be carried out concurrently. In particular, it is not reasonable to defer networking level research activities to the distant future since one cannot be expected to “build the Roman Empire in a day”. As a case in point, it is important to realize that the current downturn in the optical industry is largely due to the hype generated in the past few years during which many optical networking companies made unrealistic claims about their abilities to build and deploy all-optical and photonic networks using current technologies and which were supposedly much more cost-effective than hybrid networks or electronically switched networks. The fact remains that not only are device/component technologies not yet mature, but precious little research has been conducted on the architectures, protocols and other networking level topics. For example, architectural options including IP/X/WDM (where X can be ATM or SONET etc.) as well as protocols to automatically reconfigure a WDM network for bandwidth provisioning etc. were (and still are) being investigated in academia as research topics. And it could take a while before the knowledge base can be built and research results can be implemented in real networks.

Now that the bubble in the optical industry has burst, the industry is unlikely to invest substantially in research in optical networking (either at the networking level or at the device/component level). However, to leverage the huge progresses made in optical transmission technologies, and to ensure the continued success of optical communications, it is imperative that US government funding agencies, especially NSF, maintain a long-term vision. Lack of funding for research in this area risks the US falling behind the rest of the world in one of the most important components of a global communication infrastructure.

Progress Made in Optical Networking Research

NSF has supported in the past several research projects on hybrid (e.g., translucent) networks, wavelength-routed networks, as well as optical packet and burst switched networks. This research has mainly focused on the optical network topology, routing and wavelength assignment algorithms, survivability, and other issues. As a result of various research activities, fundamental knowledge has been accumulated over the years, which forms the basis for further advancement in the optical networking field.

For example, in the context of network architectures, it has been established that the complex IP/ATM/SONET/WDM model is neither efficient nor necessary. Instead, IP/WDM is widely regarded as the model for the next generation networks, although many issues remain open, including the exact roles of G-MPLS, (next generation) SONET, and Gigabit Ethernet. For local/access networks, simple topologies based on stars and rings have been augmented with trees, meshes and AWGs (Arrayed Waveguide)-based topologies, but much more work on optical broadband access and Metro networks needs to be done.

In addition, for wavelength-routed networks, efficient routing and wavelength assignment (RWA) algorithms that work without wavelength conversions (as mentioned earlier) or with a limited number of wavelength converters strategically placed in a network, or wavelength converters with limited conversion capabilities have been proposed and evaluated. WDM-specific issues related to optical signal degradation have also been addressed with QoS-aware RWA algorithms. Research problems related to efficient design and reconfiguration of virtual (WDM-layer) topology to support client (e.g., IP/MPLS) connections, traffic grooming algorithms to reduce the cost of electronic equipments such as SONET Add-Drop Multiplexers, distributed control and signaling algorithms (for lightpath establishment and release), WDM-layer multicast have also been studied. Tools useful for network design and performance monitoring, control and signaling software, and traffic engineering approaches such as waveband switching (WBS) are among issues that warrant further research.

Furthermore, research on network scalability and survivability issues has led to the departure from the ring-based topology used by SONET to new meshed network topologies. At the optical (WDM) layer, many mesh-based shared path and link protection and restoration schemes have been developed. Recent work on shared protection based on multiple-segments promises to improve both the recovery speed and bandwidth utilization when compared to shared path protection. Continued research effort on efficient schemes dealing with Shared Risk Link Groups (SRLGs) and multi-layer (IP/GMPLS/WDM) recovery is needed.

Compared to wavelength-routing, optical packet and optical burst switching (OBS) are relatively immature areas of research carrying a higher risk with a potentially high return as well. In recent years, interest in OBS (as a variation of OPS) has grown as is evident from scores of publications in major conferences (INFOCOM, Globecom and ICC) and journals (e.g., JSAC). Research in this area has contributed to the understanding of the advantage as well as limitations of burst assembly at the edge and bufferless switching in the core. Although current device/component technologies are not mature for OPS/OBS to be cost-effective, research in this area with long-term perspectives should be encouraged. Among the issues to be examined are the cost-performance comparison between OPS/OBS and wavelength routing, novel protocols to provide SONET-like QoS to real-time applications, advanced traffic engineering issues including efficient burst scheduling and contention resolution algorithms, and TCP performance in OBS networks

Challenges and Difficult Problems That Remain

The critical challenges in enabling the deployment of optical networking in the future may broadly be classified into the following areas (i) Access/Metro Networks (ii) Core network architectures (iii) Integrated device and network research (iv) Network management & control (v) Robust and Secure optical networking and (vi) Application-driven optical networks. Below, we highlight major open problems to be studied in each of the areas.

1. Access/Metro Networks

- The so-called “last or first mile bottleneck” currently prevents users from benefiting from the large bandwidth available in the core. Optical access technology that can provide affordable broadband connectivity to the edge of the network should be pursued as a viable alternate/complementary technology to wireless access.
- Efficient grooming, cross-connect and switching architectures for the optical Metro networks that can meet different (and time-varying) traffic granularity and QoS needs of users should be studied.
- Novel multi-service platforms for integrating access/metro networks and core networks should be developed with special emphasis on ease of deployment and management.

2. Core Network Architectures

- One of the key issues in the design of core network architectures is scalability. How to transport data across a wide area in a cost-effective manner should be a primary focus of research in this area.
- Another important issue is rapid provisioning of network bandwidth at various granularity levels including the wavelength or packet/burst level. Proposed architectures should take advantage of unique capabilities of optical devices that are either available today (e.g., wavelength add/drop for optical bypass at intermediate switches) or may become available in the future (e.g., fast switching fabrics with nanosecond switching time).
- The evolution of core optical transmission and switching technologies should also be studied taking into consideration the benefits of integrating optics and electronics where appropriate.

3. Integrating device/component and network research

- It is evident that network architectures guide device and component technology development and at the same time, device and component capabilities influence network design. Tools for the design, modeling/simulation, and evaluation of optical devices, components and networks are of particular importance.
- Multi-disciplinary initiatives and test-bed deployments can effectively accelerate research in this area, and should continue to be supported.

4. Network management & control

- Rapid bandwidth provisioning is achievable in the next-generation optical network architectures. However, further research is required to develop the control software, physical layer modulation and signaling protocols, and mechanisms for monitoring, measurement, and fault-isolation.
- The relative advantages of O/E/O vs. all-optical approaches for switching should be studied from a network management point of view.
- Standardization of control plane protocols to achieve interoperability across vendor equipments must be encouraged.

5. Robustness and security

- With increasing line speeds, protection/restoration at the optical layer is clearly more attractive than restoration at higher layers in the event of fiber cuts and/or node failures. Efficient failure-resilient and survivable network architectures and protocols that take into consideration the interactions between multiple layers to achieve fast and guaranteed recovery without incurring excessive overhead should be investigated.
- Security of data transported across optical networks should be assured through a new generation of high-speed cryptographic techniques and protocols.
- Exploiting novel physical-layer techniques (such as coding) for privacy/security should also be investigated.

6. Application-driven optical networks

- All indications are that the future will witness new applications which are even more bandwidth intensive than those available or imaginable today, and optical networks can offer unique capabilities such as higher bandwidth, scalability and security to support specialized applications such as Telemedicine and Grid Computing. Vertically integrated research effort at all layers from the application layer to the physical layer in developing both software and hardware should be encouraged.
- Both applications and optical networking technologies should be redesigned and existing protocols constantly improved to optimize for TCP/IP which is and will likely to remain the prevailing transport/routing protocols.
- On the other hand, new protocols radically different from those based on TCP/IP may be needed by a new generation of applications, and thus should also be investigated.