

MEMS: The Path to Large Optical Crossconnects

Patrick B. Chu, Shi-Sheng Lee, and Sangtae Park, Tellium, Inc.

ABSTRACT

Continuous growth in demand for optical network capacity and the sudden maturation of WDM technologies have fueled the development of long-haul optical network systems that transport tens to hundreds of wavelengths per fiber, with each wavelength modulated at 10 Gb/s or more. Micro-electromechanical systems devices are recognized to be the enabling technologies to build the next-generation cost-effective and reliable high-capacity optical crossconnects. While the promises of automatically reconfigurable networks and bit-rate-independent photonic switching are bright, the endeavor to develop a high-port-count MEMS-based OXC involves overcoming challenges in MEMS design and fabrication, optical packaging, and mirror control. Due to the interdependence of many design parameters, manufacturing tolerances, and performance requirements, careful trade-offs must be made in MEMS device design as well as system design. In this article we provide a brief overview of the market demand, various design trade-offs, and multidisciplinary system considerations for building reliable and manufacturable large MEMS-based OXCs.

INTRODUCTION

To meet the growing demand for high data bandwidth, service providers are building optical networks around the globe using the latest wavelength-division multiplexed (WDM) technologies with mesh network architecture [1]. Lightpaths between access points in a network are created using fiber links containing many wavelength channels in each fiber, where each channel or port can have a data rate of up to 2.5 or 10 Gb/s. At the edge of the networks are the clients (IP/ATM routers, optical add-drop multiplexers, etc.) that use these lightpaths as high-capacity pipes for data/voice traffic. Data rate per port is expected to continue to increase (40 Gb/s in the very near future). The number of wavelength channels (or ports) per fiber will

also continue to rise as WDM technologies mature.

For long-haul core networks, core switching is needed for two main purposes: network provisioning and restoration (Fig. 1). Provisioning occurs when new data routes have to be established or existing routes modified. A network switch should carry out reconfiguration requests over time intervals on the order of a few minutes. However, in many core networks today, provisioning for high-capacity data pipes (OC-48 — 2.5 Gb/s and OC-192 — 10 Gb/s) requires a slow manual process, taking several weeks or longer. High-capacity reconfigurable switches that can respond automatically and quickly to service requests can increase network flexibility, and thus bandwidth and profitability.

On the other hand, restoration must take place in events of network failures (e.g., an accidental cable cut). A network switch needs to reroute traffic automatically in a time interval on the order of 100 ms, thus restoring operation of the network. Traditionally, network restoration is performed primarily by digital electronic cross-connects and synchronous optical network (SONET) add-drop multiplexers, operating at a data rate of about 45–155 Mb/s. For switches in a core network handling hundreds of gigabits per second of traffic, restoration at a coarser granularity is desirable in terms of both cost and manageability. Provisioning and restoration at coarse granularities also makes sense in light of the development of high-speed service-layer equipment such as IP routers with 10 Gb/s interface and Gigabit Ethernet.

These provisioning and restoration requirements of next-generation optical networks demand innovations in switching technologies. In the following sections, a vision and technologies for next-generation optical crossconnects (OXCs) are described, with a focus on MEMS technologies as the leading choice for photonic switching. Key challenges associated with the development of MEMS-based OXCs are discussed. Finally, an outlook on MEMS-based OXC development and deployment is presented.

NEXT-GENERATION CROSSCONNECTS

An emerging vision of the next-generation crossconnects for optical networks is one that allows network reconfiguration in the optical layer (Fig. 2a): provisioning and restoration in large units (e.g., the wavelength). Since the number of wavelengths per fiber has already reached hundreds today (160 wavelengths for 10 Gb/s) and is expected to increase, the desired port counts for such OXCs are expected to be in the thousands, where scalability is a paramount concern. Such a switch must also operate in a fully nonblocking manner, where every input must be allowed to connect to every output with no restriction. In addition, insertion loss, physical size, polarization effects, and switching times are also critical considerations. Equipped with intelligent provisioning and restoration capabilities, the next-generation OXC must also meet the stringent telecommunication requirements with an operating lifetime of 20 years.

OPTICAL-LAYER SWITCH WITH AN ELECTRICAL SWITCHING CORE

An optical layer switch can be implemented using opto-electronics interfaces and high-speed electronics. Due to the advancement of state-of-the-art integrated circuit (IC) technologies, multiple vendors currently offer electronics-based optical switches, also known as O-e-O (Optical-electrical-Optical) switches, with a few hundred 2.5-Gb/s ports residing in several equipment bays. These state-of-the-art switching systems provision and mesh-restore wavelengths at a granularity of 155 Mb/s to 2.5 Gb/s. For example, Fig. 2b shows Tellium's Aurora Optical Switch™ that has 512 OC-48 (2.5 Gb/s) input ports and 512 OC-48 output ports, and can deliver a total aggregate capacity of 1.28 Tb/s. They also provision and mesh-restore

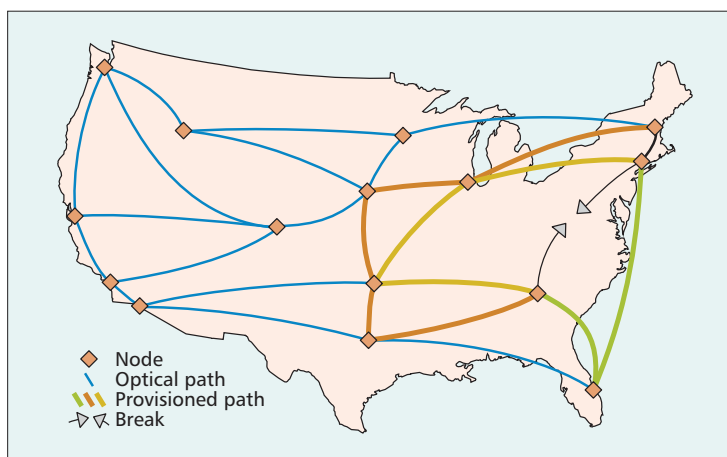


Figure 1. Illustration of data path provisioning and restoration in a core transport mesh network.

10 Gb/s wavelengths (OC-192) via inverse multiplexing down to the basic switch rate, with the capability of grooming such substrate signals within a given 10 Gb/s pipe. Intelligence of this switch allows dynamic and automatic provisioning and protection as well as in-service system upgrades. Based on multiple stages of Clos structures [1], these switches are also scalable to thousands of switching ports.

OXCs WITH MEMS-BASED OPTICAL SWITCHING CORE

OXCs with electrical switching cores like the Aurora Optical Switch will continue to be deployed and remain in service for quite some time. Higher-speed and higher-capacity electronics switches are expected to reach the market in the near future as IC technology advancement

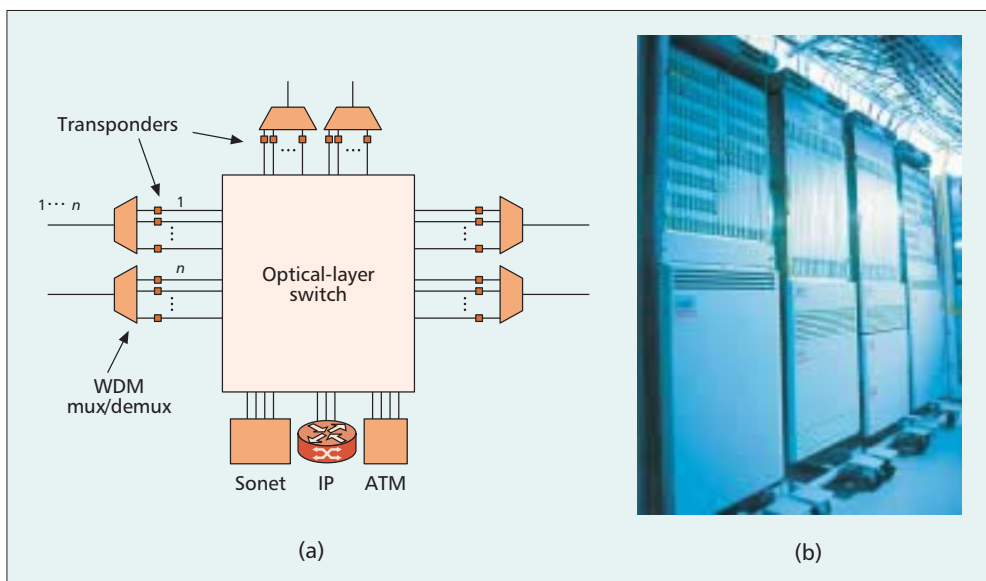



Figure 2. a) Illustration of an optical-layer switch connected to DWDM transport systems and client equipment; b) Tellium's Aurora Optical Switch™ with 512 OC-48 (2.5 Gb/s) input ports and 512 OC-48 output ports, 1.28 Tb/s of aggregate switching capacity deployed, and carrying commercial traffic today.



■ **Figure 3.** a) Illustration of a 2D switching architecture; b) 2D $N \times N$ switches first demonstrated by AT&T [8].

continues. On the other hand, the possibilities of improved scalability, footprint, manageability, and cost continue to fuel the quest for technological solutions beyond the proven state of the art. A new concept that has arisen is an all-optical OXC: an optical-layer switch with an optical switching core. All-optical switches are also known as *O-o-O* (Optical-optical-Optical) switches, which can be realized using arrays of MEMS-fabricated micro-mirrors.

MEMS for Photonic Switching — MEMS technology enables the fabrication of actuated mechanical structures with fine precision that are barely visible to the human eye. MEMS devices are by nature compact and consume low power. A batch fabrication process allows high-volume production of low-cost devices, where hundreds or thousands of devices can be built on a single silicon wafer. While the MEMS field is young compared to traditional semiconductor electronics, MEMS technology is based on fabrication technology fundamental to IC fabrication and many mature engineering disciplines such as mechanics, electromagnetics, and material science. Applied research in MEMS over the past two decades has led to numerous successful commercial devices, including valves and pressure sensors for automotive and medical applications, accelerometers, and angular rate sensors for airbags, toys, and instrumentation on land, at sea, in air, and in space. On the other hand, technological wonders such as injectable micromachines performing heart surgery inside the human body will remain fantasies of fiction writers for many decades to come.

Optical MEMS, nevertheless, is a promising technology to meet the optical switching need for large-port-count high-capacity OXCs. Within the last decade, the realization that tiny micro-machined structures can steer light by generating small tilting motions has opened doors to many exciting applications of MEMS in photonic switching [2–4]. Current (nonelectronics) competing technologies for building are thermal bubble switches, which make use of total internal reflection and index-matched fluid, and waveguide-based switches, which make use of inter-

ferometric effects of light in planar waveguides. Potential benefits of an all-optical MEMS-based OXC include scalability, low loss, short switching time, low power consumption, low crosstalk and polarization effects, and independence of wavelength and bit rate. Therefore, MEMS has become the leading choice of technology for building large all-optical OXCs.

The most notable commercial MEMS optical devices to date are Texas Instruments' Digital Mirror Devices (DMD) [5], which have found applications in consumer visual display and projectors. While different MEMS-based solutions for critical transmission applications such as gain equalization [6] and dispersion compensation [7] are under investigation, add-drop multiplexers and small protection switches are among MEMS-based optical products that are slowly reaching the market. In recent news, small optical switch products have been announced to pass rigorous Telcordia telecommunications specifications, beginning to cast away healthy doubts about the long-term reliability of MEMS devices. Large MEMS-based OXCs as fully qualified products are expected to be a reality in the near future.

Two-Dimensional MEMS Switches — The OXCs of main interest are fully nonblocking optical switches with N input and N output ports. Two architectures for MEMS-based OXCs have emerged. In the first architecture, often known as 2D switching (Fig. 3) [2, 8, 9], a square array of $N \times N$ mirrors is used to couple light from a linear array of N fibers on one side of the square to a second linear array of N fibers on an adjacent side of the square. The (i, j) mirror is raised up to direct light from the i th input fiber to the j th output fiber. Mirror control for these 2D switches is binary and thus straightforward, but the trade-off of this simplicity is optical loss. While the path length grows linearly with N , the number of ports, the optical loss also grows rapidly due to the Gaussian nature of light. Therefore, 2D architectures are found to be impractical beyond 32 input and 32 output ports. While multiple stages of 32×32 switches can theoretically form a 1000-port switch, high optical losses also make such an implementation impractical.




Figure 4. a) Illustration of 3D switching architecture; b) illustration of beam steering using a two-axis gimbaled mirror; c) fabricated MEMS gimbaled mirror array.

Three-Dimensional MEMS Switches — In the 2D case, all the light beams in the switch reside on the same plane, resulting in unacceptably high loss for large port counts. The second architecture (Fig. 4a), known as 3D switching [10–12], makes use of the three-dimensional space as an interconnection region, allowing scaling far beyond 32 ports with acceptable optical losses (< 10 dB). In this architecture, there is a dedicated movable mirror for each input and each output port. Each mirror must now operate in an analog, rather than binary, mode, tilting freely about two axes (Fig. 4b, c). This elegant architecture offers the virtue that the optical path length now scales only as \sqrt{N} instead of N , so port counts of several thousand are achievable with losses below 10 dB. This 3D optical architecture clearly presents real hope for developing a scalable large-port-count OXC.

THE PATH TO A MEMS-BASED OPTICAL CROSSCONNECT

MEMS-based OXCs are no doubt feasible in concept. Substantial challenges must be overcome for any switch design; these challenges include MEMS mirror manufacturing, optomechanical packaging, and mirror control. Many aspects of these three challenges are interdependent. Complex trade-offs must be weighed in designing a MEMS-based OXC.

MEMS DESIGN AND FABRICATION

Components of a large MEMS-based OXC include thousands of actuated mirrors, lenses, collimators, and fiber arrays. With no doubt MEMS mirrors, the key active element in the optical system, are the most critical technology for large OXCs.

MEMS Design — A two-axis actuated tilting mirror can be divided into three elements: the mirror, the springs as the mechanical support, and the actuator, all of which determine important OXC system parameters. Examples of these parameters include maximum port count (dependent on the mirror tilt angle), switch settling time (dependent on the mirror response time), insertion loss (dependent on the mirror size, reflectivity, and maximum tilt angle), and power dissipation (dependent on power required for mirror actuation and control). For a 1000-port switch, each mirror may require a diameter on the order of 1 mm, with mirror radius of curvature (ROC) greater than a few tens of centimeters. Reflectivity of each mirror is desired to be above 97 percent. The tilt angle requirement ranges from a few degrees to $\pm 10^\circ$ depending on the optical train design of the OXC.

The challenges in MEMS design come from the different trade-offs between desired properties of the MEMS device. As an example, the supporting springs for the mirrors must have sufficient stiffness to meet the mirror response time and vibration immunity requirement. But the upper bound of the spring stiffness is determined by the desired maximum tilt angle and the actuator's maximum force or torque output (as well as the switch power budget). Magnetic actuation and electrostatic actuation are two viable choices for mirror positioning. Magnetic actuation offers the benefit of large bidirectional (attractive and repulsive) linear force output but requires a complex fabrication process and electromagnetic shielding. Electrostatic actuation is the preferred method mainly because of the relative ease of fabrication and integration. However, to achieve large tilt angle using a stiff spring, the trade-offs include high actuation voltages (on the order of 50–200 V) and nonlinear torque output.

For a typical Z-configuration 1000-port switch, coupling losses between the input and output fibers can be computed using Gaussian beam propagation methods. Component fabrication tolerances and packaging tolerances can also be estimated.




Figure 5. a) Top view of a MEMS mirror; illustration of an SOI-based electrostatic MEMS mirror; b) before; and c) after structural release of the gimbaled mirror.

A particular challenge for MEMS mirror design is to maximize ROC. A stable metal coating such as of gold, along with necessary additional metal adhesion and diffusion barrier layers, is often used as a reflective surface. These metal coatings can create an undesirable temperature-dependent mirror curvature due to intrinsic stress of the metal layers and the difference in thermal expansion coefficients of the metal coating layers and the bulk mirror made of a different material. This problem is especially severe if the metal coating is applied only to one side of the bulk mirror. A thick mirror can best counteract curvature from stress induced by metal coating on the mirror. Unfortunately, large mass leads to slow mirror response time and high sensitivity to stochastic vibration.

MEMS Fabrication Choices — In principle, the bulk mirror can be made of any material as long as reliability, reflectivity, and optical flatness requirements are met. Single-crystal silicon (SCS), commonly used in MEMS, is recognized to be the most suitable choice over polysilicon or electroplated metal due to low intrinsic stress and excellent surface smoothness. The choice of material for the mirror springs is arguably even more important because the mirror springs will constantly be twisted and bent. Superior mechanical characteristics make SCS the best candidate for the mirror springs. Alternative materials such as polysilicon and metal are poor substitutes because of potential stress, hysteresis, and fatigue problems. In most cases, the same material is chosen for both the bulk mirror and the springs in order to yield a straightforward fabrication process.

A plethora of fabrication processes can be used to create two-axis actuated SCS mirrors or mirror arrays [11, 13, 14]. Besides typical lithography, deposition, and etching procedures, necessary fabrication steps may include deep reactive ion etches (DRIE), silicon wafer bonding, and chemical mechanical polishing (CMP) [5]. Silicon-on-insulator (SOI) wafers are a convenient starting material to create SCS bulk mirrors with uniform thickness and low intrinsic stress (Fig. 5), but these wafers are unfortunately expensive with few supply vendors today. Applying clever silicon etching and wafer bonding techniques to cost-effective [100]-type silicon wafers may also yield mirrors with sufficiently low mass and large ROC. The pri-

mary differentiating factor between these MEMS mirror processes is device performance characterized by mirror size, flatness, reflectivity, maximum mirror tilt angle, and ease of mirror control. Material supply availability, length of fabrication cycles, and equipment bottlenecks play important roles in shortening product development cycle and time to market. Ease of circuit integration, achievable mirror array fill-factor, mirror array size, and manufacturing yield may also influence the overall switch fabric design. Arguably, a fabrication process that enables monolithic integration of electronics with MEMS [14] may lead to MEMS mirrors with the greatest functionality and the highest performance.

OPTICAL PACKAGING

The optical system as shown in Fig. 4a requires thousands of micro-mirrors, lenses, and fibers aligned to each other with tolerances on the order of microns and hundreds of micro-radians. This multi-element body must endure thermal cycles, shock, and vibration during shipping and operation, which may lead to short-term and long-term mechanical drift in packaging. Obviously, tolerance of various pointing errors and misalignment errors depends on the robustness of the optical architecture design. In addition, these thousands of optical components must be carefully and compactly packaged with all the necessary control electronics in order to meet the additional space constraints and front panel accessibility requirements of telecommunications equipment.

For a typical Z-configuration 1000-port switch like Fig. 4b, coupling losses between the input and output fibers can be computed using Gaussian beam propagation methods. Component fabrication tolerances and packaging tolerances can also be estimated [4]. For example, ± 1 percent of focal variation in a single port lens in a lens array could account for up to 1 dB of optical loss. $\pm 2 \mu\text{m}$ of relative position error in a fiber array can also lead to similar losses. One method to facilitate packaging is to make use of large fiber bundles, lenslet arrays, and monolithic dies with thousands of mirrors. The number of optical elements in the system may then be reduced to half a dozen or so. However, fabrication and packaging of such large fiber bundles, lenslet arrays, and MEMS mirror dies poses formidable challenges of their own.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.