Lightwave micromachines for optical networks:
Vast promise amid vaster promises

E. L. Goldstein, L. Y. Lin, and J. A. Walker

Tellium, 2 Crescent Place, Oceanport, NJ 07757 USA

(732) 923-4256 (732) 923-9804 (fax) evango@tellium.com

Abstract. As data-networking interfaces edge towards the per-wavelength line rate, while core-transport line rates steadily rise, there is an emerging need for reconfigurable networking at the granularity of an individual wavelength, currently 2.5 to 10 Gb/s. This need places severe strains on an already-war-weary arsenal of conventional opto-electronics, and thus offers striking opportunities for lightwave micromachines, or micro-electro-mechanical systems (MEMS), that were implausible only a year or two ago.

There have been two near-immediate results. Micromachines have emerged as a genuinely disruptive technology whose communications promise has few precedents. Concurrently, widespread hyperventilation about the promise and ubiquity of micromachines for optical networking (cf. Fig. 1) has itself reached a level with few recent precedents.

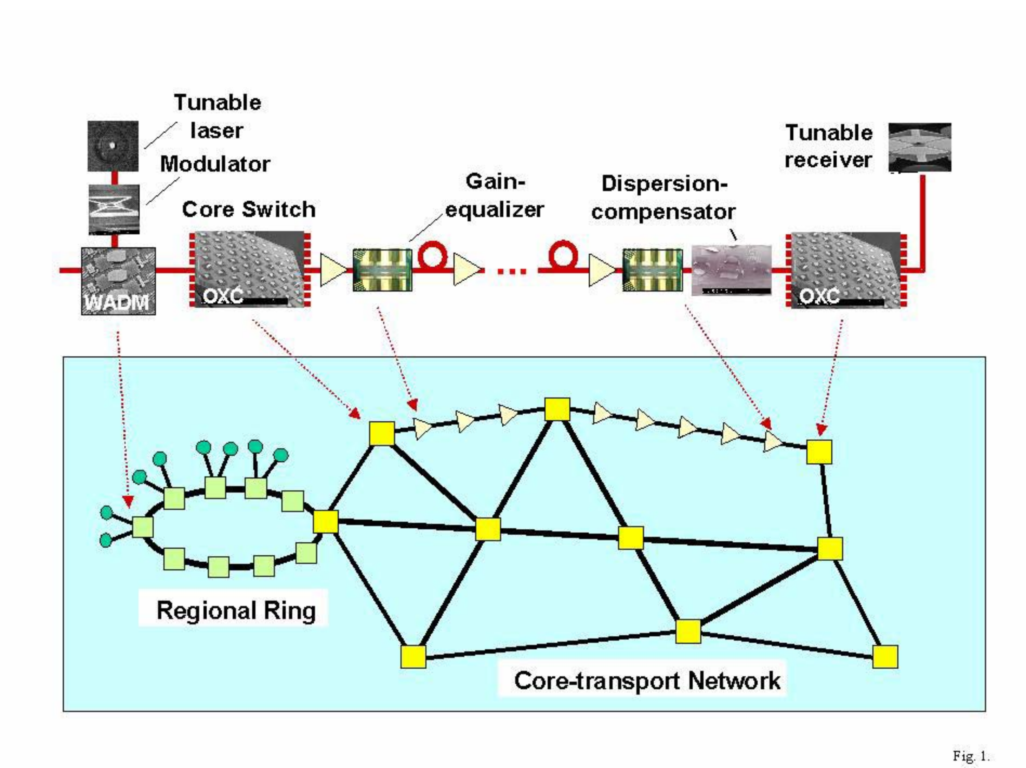

Fig. 1: *Optical-layer networking with lightwave micromachines: the cheerful vision.*

The resulting tension seems almost certain to produce a mixed outcome in which the genuinely sweeping impact of lightwave micromachines over the long-term is tempered by some disappointment with the timescale over which this impact is felt. In long-haul transmission-system impairment-mitigation, where the arsenal of conventional opto-electronics offers few strong alternatives, MEMS represent serious candidates for satisfying the network's rapidly mounting needs for tunable equalization of gain, chromatic dispersion, and polarization-mode-dispersion, and for reconfigurable wavelength-add/drop. Here, some near-term impact is likely. On the other hand, in large core-network switching systems, where micromachines face far deeper technical and competitive challenges, they nonetheless offer the long-term promise of raw, aggregate switch capacity that is unmatched by other technologies. Still, notwithstanding a rising stream of tour-de-force concept-demonstrators, the serious deployment of MEMS as core-network switches will almost certainly occur over timescales that disappoint current overheated expectations. We here summarize the genuinely disruptive promise of lightwave micromachines in core-transport optical networking, as well as the substantial challenges they have yet to overcome.

1. Lightwave micromachines for transmission

The globe's core long-haul communications networks are now caught in the midst of an onslaught from three diverse sources. Aggregate demand is swiftly rising. This demand is rapidly transforming from time-multiplexed circuit-switched traffic to packetized data, with swift convergence to the Internet Protocol. Finally, bit rates of individual services are rapidly rising, with 2.5 and 10 Gb/s private-line services abruptly mushrooming, owing to the sudden emergence of IP router interfaces at those rates. These three changes have over the past five years collectively driven the most fundamental physical transformation that optical-fiber systems have yet experienced: transmission fibers have gone from simple, unamplified, single-channel 2.5-Gb/s pipes to optically amplified, Tb/s systems transporting hundreds of wavelength-division-multiplexed (WDM) light signals. It is chiefly this sudden proliferation of wavelengths in WDM transport systems that is just now beginning to drive an increasingly acute need for optical subsystems that can switch, wiggle, and adjust; thus, the abrupt recent epiphany over MEMS.

In 1995, the inaugural year for commercial WDM transmission-system deployment, a total of eight 2.5 Gb/s channels were painfully coaxed through recalcitrant amplified systems, achieving a capacity of 20 Gb/s per fiber. Apart from optical gain, no equalization technologies at all were used to combat the effects of transmission impairments. The art has since accelerated with devastating speed. The year 2001 will witness the commercialization of 2.4 Tb/s systems transporting 240 10-Gb/s channels. It will also see the first tentative deployments of WDM systems operating at 40 Gb/s per wavelength.

As a result of this acceleration, total optical power levels traversing fiber links are precipitously rising, while spectral channel spacing and bit period simultaneously shrink. This in turn means that optical-amplifier noise, fiber nonlinearity, and dispersion become major players requiring meticulous transmission-engineering, thus creating rich opportunities for any technology that can mitigate accumulating optical performance-impairments. Micromachines provide one such particularly promising technology.

Gain-equalization. One of the earliest and strongest opportunities for MEMS is in the mitigation of spectral gain nonuniformities. These result from the erbium-doped fiber amplifiers that have made long-haul WDM transport feasible by . Erbium amplifiers generate gain spectra of the form

$$G(\lambda) \sim \exp\int_0^L \left[\sigma_e(\lambda) n_2(z) - \sigma_a(\lambda) n_1(z)\right] dz),$$

where σ_e and σ_a , the erbium emission and absorption cross-sections, are weighted by the excited and ground-state populations n_2 and n_1. The resulting ragged set of physically feasible gain shapes, generated by the lumpy cross-sections of Er atoms in alumino-silicate glass, is shown in Fig. 2 [1].

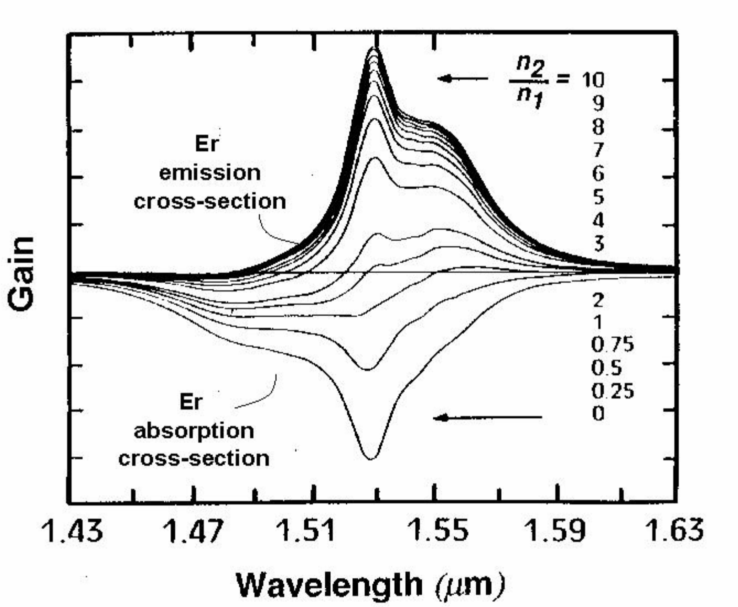

Fig. 2. *Erbium amplifier gain spectra—origin of the tight parameter-coupling in core-transport networks. Curves depict the family of feasible erbium amplifier gain spectra, exhibiting their sensitive dependence on population-inversion.* n_2*: excited-state population.* n_1*: ground-state population.*

Owing to saturation effects, the above gain expression implies that amplified WDM links are by their nature extremely strongly coupled systems, whose end-to-end transmission spectra are sensitively determined by essentially all major parameters in the system—channel-count, pump and signal power-levels, inter-amplifier loss—all of which impact the populations n_2 and n_1, and thus the gain shape. It is this strong coupling that generates large variations in end-to-end transmission spectra that are difficult to predict and control [1]. Thus ultimately requires dynamically adjustable gain-equalization.

This need is exacerbated both in systems that incorporate wavelength-add/drop and in wideband (~80 nm) systems, where parasitic Raman gain generates large channel-loading-dependent gain variations [2],[3]. Since high-channel-count systems inevitably require both wavelength-add/drop, to avoid stranding capacity, and large optical bandwidths, due to four-wave-mixing and optical filter-resolution constraints, such systems also inevitably require dynamic gain-equalization. This set of transmission considerations in practice becomes acute as channel counts rise from the tens to the hundreds—a transition that is underway right now in deployed systems. Thus, the need for dynamic gain-equalization in WDM transport systems is now becoming urgent.

As the need for dynamic gain-equalization has accelerated, technologies ranging from acousto-optics to liquid crystals to waveguide grating routers and micromachines have begun to show promise for addressing it. Among MEMS approaches, the most successful to date is almost certainly the continuous micromechanical attenuator array based on the well-known MARS structure [4]. This is shown in Fig. 3. The silicon-nitride membrane plus air gap serve as a spatially variable, tunable multi-layer dielectric mirror. When the incident signal is spectrally dispersed along the axis of the device defined by an array of strip electrodes, as in Fig. 3b, one obtains a simple and compact tunable spectral shaper. Such devices have been shown capable of generating a rich variety of filter shapes that quite effectively flatten out the lumpy gain spectra of erbium amplifiers, as indicated in Fig. 3 c and d.

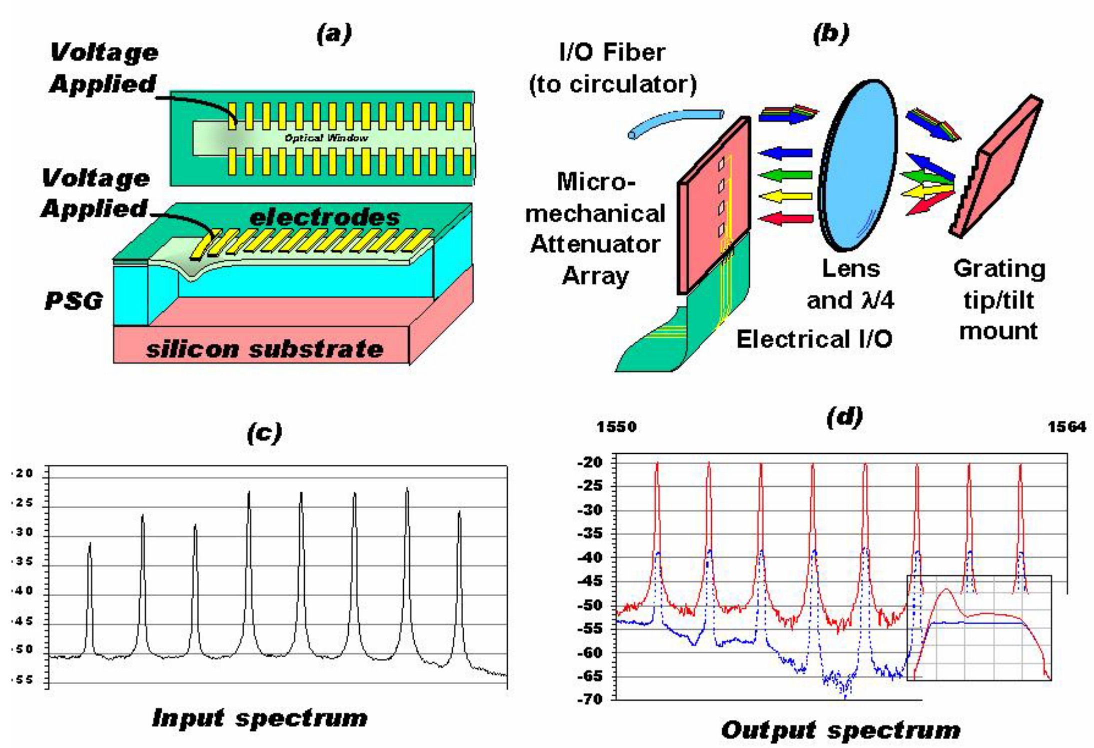

Fig. 3. *MEMS-based dynamic gain-equalizer. (a) Continuous-membrane attenuator array. (b) MEMS attenuator array embedded in a free-space wavelength multiplexer to yield a dynamically tunable spectral shaper. (c) Input signal constellation after erbium amplifier. (d) Output signal constellation after gain-equalizer (dashed line) and reamplification (solid line). Inset shows flattened spontaneous-emission spectrum. Figures courtesy of Lucent Technologies.*

Dispersion-compensation. As per-channel bit rates rise, WDM transport systems require tunable equalization not only of gain shapes, but also of various forms of dispersion, which causes pulse-distortion and thus inter-symbol interference. At about 40 Gb/s, the need arises for dynamically tunable compensation of the transport system's chromatic dispersion. Here too MEMS have already begun to show substantial promise. MARS-type structures, when employed as tunable mirrors in Fabry-Perot cavities, are effective means of subjecting the various spectral components of an incident signal to adjustable Fabry-Perot lifetimes, thus producing dynamically tunable dispersion values [5]. These can then be used to algebraically cancel chromatic dispersion imposed by fiber, multiplexers, and other transmission-system components.

Other forms of dispersion, too, require dynamic compensation. Due to departures from precise radial symmetry in the cores of real fabricated and deployed optical fibers, the two orthogonal polarizations in fact travel with different group velocities, resulting in distortion due to polarization-mode-dispersion (PMD). In digital systems, one needs to hold the time-averaged differential delay between these two polarizations,

$$\langle \Delta \tau \rangle = \sqrt{\sum_{k=1}^{M} D_{PMD}(k)^2 * L(k)}$$

to values smaller than about 0.1 *T*, where *T* is the bit period, $D_{PMD}(k)$ is the PMD parameter of the kth fiber segment, and L(k) is that segment's length. Since the PMD parameter $D_{PMD} \sim 1$ ps/(km)$^{1/2}$ for typical embedded fiber, this implies maximum transmission lengths of about 1600 km at 2.5 Gb/s, 100 km at 10 Gb/s, and 7 km at 40 Gb/s. Thus, while PMD is in practice not a factor at 2.5 Gb/s, its compensation becomes important at 10 Gb/s and rapidly grows critical at higher rates. Because the orientations of the fiber's slow and fast axes vary in time, compensation must be dynamic. Moreover, because the compensation must be done on a channel-by-channel basis, PMD-compensators, if they are to compete with the alternative of simply replacing older installed PMD-afflicted fiber, must be compact and cheap. This again presents a natural opportunity for micromachines.

No mature MEM-based PMD-compensating technology has yet been reported even in the research literature. However, the kernel of such a compensator would be a polarization-controller which, in combination with an appropriate birefringent element, provides first-order PMD-compensation. The basic building block of this, a polarization-rotator, has been implemented via the MEMS-based device of Fig. 4 [6]. This consists simply of a Mach-Zehnder interferometer fabricated via surface-micromachining in polysilicon, using the Chronos MUMPS process [7]. Since the micromachined poly plates can be held in precise position with resolution far better than an optical wavelength, the interferometer is capable of effecting precise, stable polarization-rotation via adjustment of the phase difference between |TE> and |TM> waves. This is achieved by electrostatically adjusting the position of the phase-shifting mirror shown in the SEM photograph of Fig. 5a. Applying a sinusoidal voltage to this phase-shifter should in principle generate output polarization states lying in circle on the Poincare sphere (Fig. 4). As seen in Fig. 5b, this is precisely what occurs, with operation that is stable and nearly ideal, save for a small amount of phase-to-amplitude coupling at the phase-shifter mirror, evident in slight departures from circular trajectories.

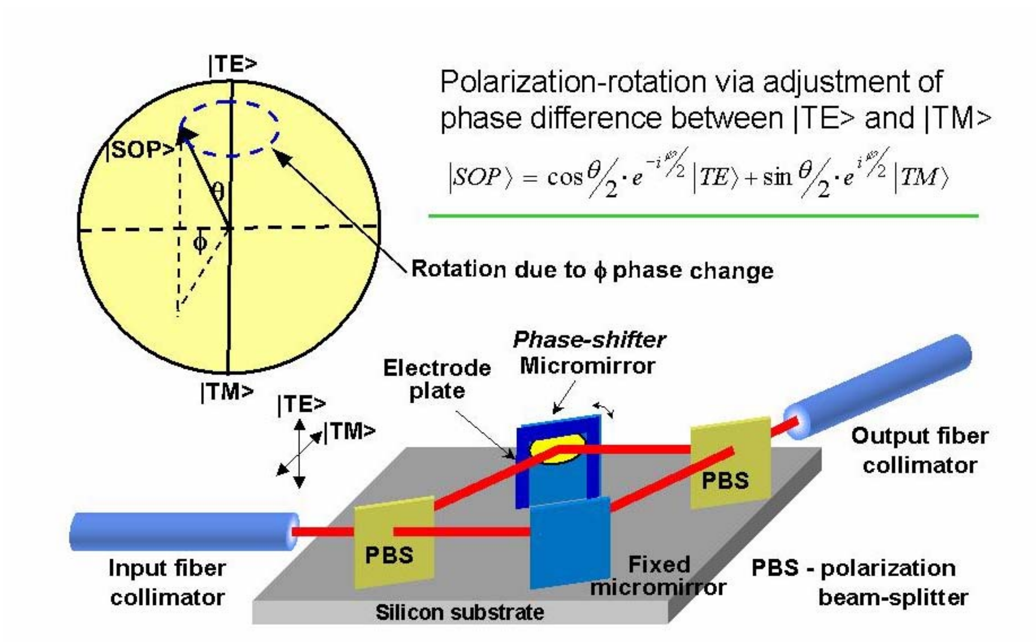

Fig. 4. *Operating principle of micromachined polarization-rotator.*

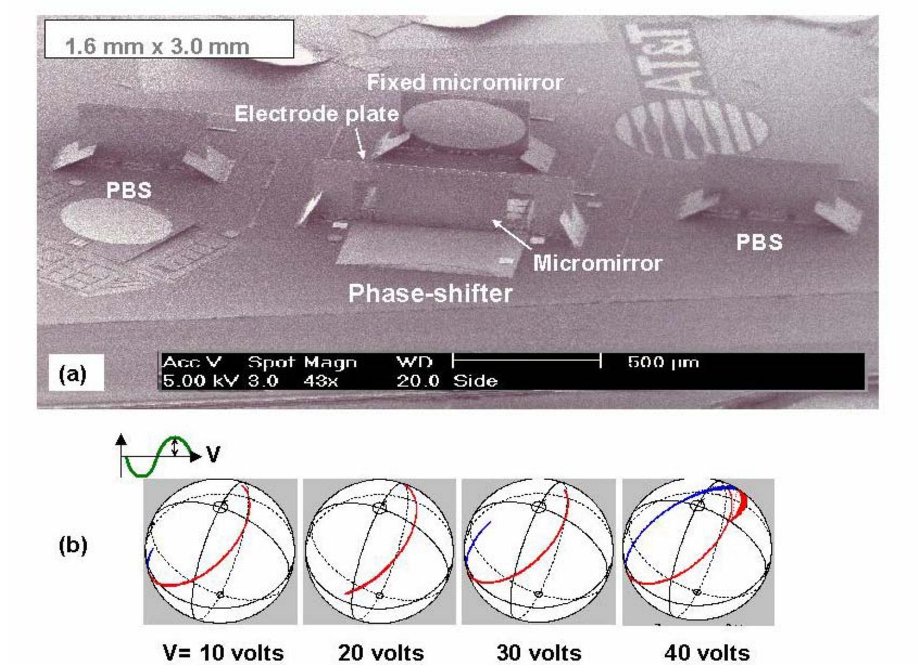

Fig. 5. *(a) SEM photograph of surface-micromachined polysilicon polarization-rotator. (b) Polarization-rotation trajectories on the Poincare sphere for sinusoidal applied voltages of various amplitude.*

In both of these applications—as dynamic gain-equalizers and as dispersion-compensators—micromachines have with remarkable speed begun to move out of the domain of the research lab and into contention as candidates for deployment in WDM transmission links. MEMS-based devices of similar complexity, generally fabricated via surface-micromachining of polysilicon, have also emerged as candidate means of reconfigurably adding and dropping wavelengths in WDM transmission links [8], [9]—a capability that becomes critical, in order to avoid stranding capacity, as wavelength channel-counts grow into the hundreds.

2. Lightwave micromachines for switching

The discussion thus far has been restricted to micromachines for transmission. And yet, with the maturation of the sources, amplifiers, compensators, filters, and high-speed electronics that have resulted this year in the first commercial terabit-per-second WDM links, core-networking challenges are swiftly shifting from transmission to switching. It is no longer problematic to transport monstrous information capacity from point to point; the difficulty is in then parting it out and disposing of it.

There are two sources of acute switching pressure in the core long-haul network. The first is *restoration*—traffic must be automatically rerouted in the event of failures over time intervals on the order of ~100 ms. The second is *provisioning*—new circuits must be established in response to service-layer requests over time-intervals on the order of a few minutes. Both functions are currently carried out using electronic switches and add/drop-multiplexers that operate typically on signals at 45-to-155-Mb/s data rates—an arrangement that becomes both unmanageable and unaffordable in core networks, just now emerging, whose larger nodes transport hundreds of Gb/s of traffic. Moreover, as earlier mentioned, service-layer vehicles—IP routers—are now sprouting interfaces, at 2.5 to 10 Gb/s, that are well-matched to the per-wavelength bit rates of transport systems. Thus, there is an emerging need to provision and restore traffic at much coarser granularity, at or approaching the wavelength. What is needed, then, is a switch scalable to thousands of ports, with each port transporting 2.5-to-10-Gb/s signals, scalable to 40

Gb/s, and with switching times <10 ms. This challenge greatly outstrips the capabilities of embedded networking technology. The result has been a frenzied drive, over the past two years, to develop coarse-granularity circuit switches with capacities that stretch into the terabits per second and beyond.

Among optical technologies, the most promising means of meeting this challenge is with micromachines. Moreover, the only promising MEMS-based approaches are those that employ free-space propagation, since waveguide structures, in practice confined to shallow bends in the plane, have thus far proven incapable of achieving the dense interconnection of optical signal paths that is needed to accomplish high-port-count strictly nonblocking switching.

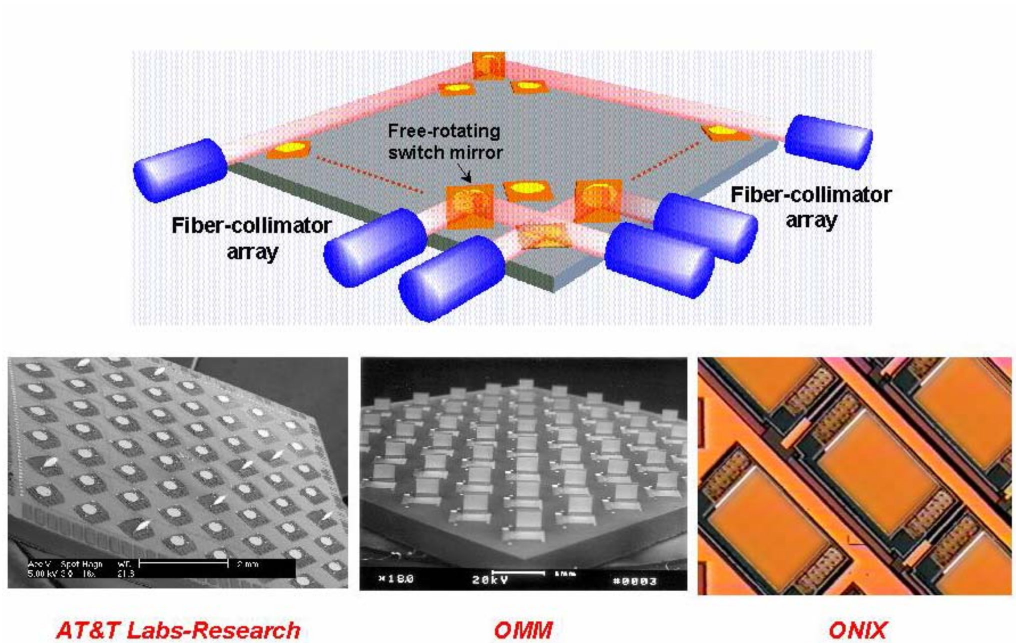

Fig. 6. *Two-dimensional cross-bar MEMS switches*.

There are two emerging approaches to building such switches using lightwave micromachines. The first is the two-dimensional cross-bar switch of Fig. 6. Here the (i, j) mirror of a square array, when raised, serves to couple the i^{th} port of the input fiber array with the j^{th} port of the output [10], [11], [12]. The advantage of this approach is that the mirror-positioning is binary—each is either up or down, but not in between. This vastly simplifies control. On the other hand, the disadvantage is loss. In such a switch, for a given Gaussian beam diameter w_0, and a fixed mirror diameter R, the optical path length grows linearly with the number of ports. Switch loss due to clipping of Gaussian beams thus grows rapidly with port count, as indicated in Fig. 7. While switch loss in principle declines monotonically as beam width is increased, beyond 32 ports the MEMS die size rapidly grows impractical, as indicated in Fig. 7. This in practice limits two-dimensional cross-bar switches to about 32 input and 32 output ports.

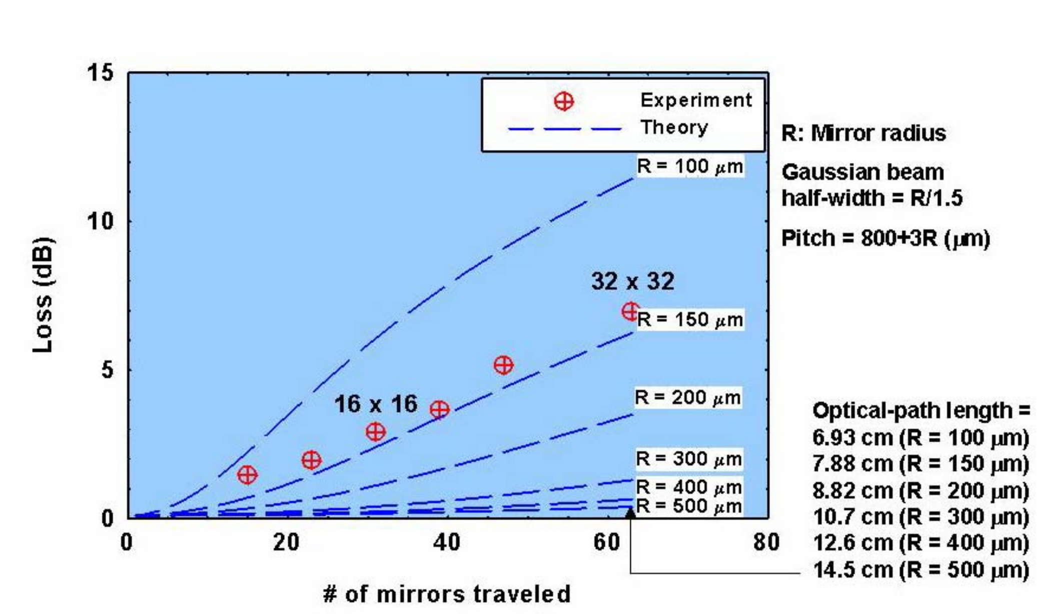

Fig. 7. *Scaling of loss, due to clipping of Gaussian beams, with port count in two-dimensional MEMS cross-bar switches.*

To scale beyond this limit one must free the beams from the constraint of lying in a plane and utilize the full three dimensions as an interconnection region. This is achieved in the three-dimensional steered-beam configuration of Fig. 8. Here a given input or output port is associated with a unique two-axis steerable mirror [13], [14], [15]. The great virtue of this approach is that for a given beam width, the optical path length scales only as $\sqrt{N}$, where N is the number of input ports. Port counts of several thousand are feasible with losses well below 10 dB. The great disadvantage of three-dimensional steered beams, however, is that they are systems containing thousands of elements—microlenses, fibers, micromirrors—each of which is constrained by truly oppressive alignment tolerances on the order of microns and hundreds of microradians. These are consequently systems that require exquisite engineering of their mirror-control servo algorithms, their fiber bundles, their lens arrays, their micromirror mechanical resonances, optical flatness, and electrical properties, and their thermal and opto-mechanical packaging. These imposing technical difficulties continue to make high yield at low loss problematic in practice, but their solutions are steadily becoming better understood.

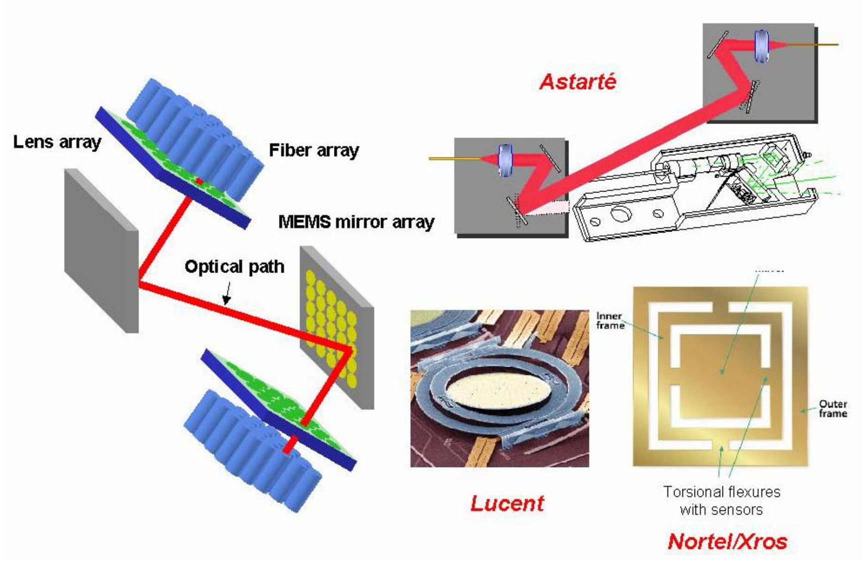

Fig. 8. *Three-dimensional analog steered-beam MEMS switches.*

3. Challenges facing the vision

Still, notwithstanding its vast promise in the long-term, substantial challenges lie between the cheerful vision articulated above and its readiness for deployment in reliable core-transport networks. The most prominent of these concerns issues of device reliability, ease of manufacture, and packaging, all of which are almost completely unexplored in the published literature to date. And yet we are proposing to deploy, in the heartbeat positions of the core-transport network, through which all high-speed links pass, network elements whose underlying technology has just been invented over the last 1-2 years. This suggests that the widespread embrace of MEMS switching systems will likely be gated by the emergence of a compelling reliability story—something that, despite growing efforts, has not yet begun to happen.

The second major challenge comes from rapid progress in electronic switching systems. Three years ago, electronic switches offering just eight 2.5-Gb/s ports seemed well beyond reach. By contrast, the current state of the art, being advanced by multiple vendors, offers switching systems with 512 2.5-Gb/s ports, for a total capacity well in excess of 1 Tb/s, residing in four equipment bays. These systems provision and mesh-restore wavelengths at a granularity of 155 Mb/s to 2.5 Gb/s. They also provision and mesh-restore 10-Gb/s wavelengths via inverse multiplexing down to the basic switch rate, with the capability of grooming such sub-rate signals within a given 10-Gb/s pipe.

Moreover, electronic technology does not yet show signs of diminishing progress. The aggregate information-handling capacity of a monolithic electronic switch chip continues to double roughly each year, as summarized in Fig. 9, with no signs yet of tapering off. 128 x 128 chips operating at 2.5 Gb/s have just become available in four distinct device technologies. When utilized in conventional three-stage switch structures, these permit the construction of strictly nonblocking 8000 x 8000 2.5-Gb/s switches, or of 2000 x 2000 10-Gb/s switches. Given this scaling, such switches will maintain a comfortable margin over the needs of even the largest core-switching nodes for some years to come. The effect of this appears to be a pushing-off of the era of the optical switch fabric by at least 1-2 years or more.

	Electronic Technology				
Monolithic Switch Port-count	GaAs	SiBipolar	SiGe	BiCMOS	CMOS
16 x 16	4Q97	1Q99	2Q99	-	
32 x 32	2Q99	-	4Q99	-	
64 x 64	4Q99	-	2Q00	4Q99	
128 x 128	1Q01	-	1Q01	4Q00	4Q00

Fig. 9. *State of the art in monolithic electronic strictly nonblocking switch chips. Each port operates at 2.5 Gb/s.*

The final challenge comes from the need to incorporate intelligence into the core-network switch—a need that owes to the fact, long-lamented by network-operators, that it costs more to operate a core-transport network than it does to buy it in the first place. Because of this, network operators continue to demand a rich variety of management features from their switches, ranging from performance-monitoring to connection-verification to the ability, by reading and writing signal overhead bytes, to automatically uncover the identity of neighboring switches and construct network topology maps. In the long run, significant transparent-optical-switch innovation in these areas seems likely; but over the next several years, the practical approach to these challenges will be to seek raw capacity from optical switch fabrics, while recovering sophisticated management and network-monitoring functions by utilizing optoelectronic interfaces.

Summary. It is at root the sudden proliferation of wavelengths in WDM systems that has created a sudden need for compact, high-density subsystems that can switch, tune, and adjust. In long-haul transmission systems, lightwave micromachines represent relatively near-term candidates for satisfying needs—just now rising to the level of urgency—both for tunable equalization of gain, chromatic dispersion, and PMD, and for implementing reconfigurable wavelength-add/drop. In core-network switching systems, on the other hand, where micromachines face substantially deeper challenges, they nonetheless offer promise, in the long-term, of aggregate switch capacity that is unmatched by other technologies. Since the globe's appetite for transmitted bits as yet shows no sign of sating, making use of the eventual raw capacity of MEMS core switches will probably not, after all, require the imagination of the poet; it will likely, however, require some measure of his patience.

FIGURE CAPTIONS

REFERENCES

[1] C. R. Giles and E. Desurvire, "Modeling erbium-doped fiber amplifiers," *IEEE Journal of Lightwave Technology*, 9 (2), 1991, pp. 271-283.

[2] A. R. Chraplyvy and R. W. Tkach, "What is the actual capacity of single-mode fibers in lightwave systems?" *IEEE Photonics Technology Letters*, 1993, pp. 666-68.

[3] M. Eiselt, "The impact of nonlinear fiber effects on fiber choice for ultimate transmission capacity," *Optical Fiber Communications Conference*, Baltimore, 2000, paper TuE1.

[4] J. E. Ford and J. A. Walker, "Dynamic spectral power equalization using micro-mechanics," *IEEE Photon. Technol. Lett*. 10, October 1998, pp. 1440-42.

[5] C. K. Madsen, J. A. Walker, J. E. Ford, K. Goossen, G. Lenz, "A tunable dispersion compensating MARS all-pass filter," *European Conference on Optical Communications*, Nice, France, 1999.

[6] L. Y. Lin, E. L. Goldstein, N. J. Frigo, and R. W. Tkach, "Micromachined polarization-state controller and its application to polarization-mode-dispersion compensation," *Optical Fiber Communications Conference*, Baltimore, March 2000.

[7] JDS Cronos Integrated Microsystems, *http://www.memsrus.com*.

[8] J. E. Ford, V. Aksyuk, D. J. Bishop, and J. A. Walker, "Wavelength add-drop switching using tilting micromirrors," *J. Lightwave Technol*., 17, pp. 904-911, (1999).

[9] C. Pu, L. Y. Lin, E. L. Goldstein, and R. W. Tkach, "Micro-Machined Optical Add/Drop Multiplexer with Client Configurability," *Fourth International Conference on Micro-Opto-Electro-Mechanical Systems,* Kauai, Hawaii, August 2000.

[10] L. Y. Lin, E. L. Goldstein, and R. W. Tkach, "Free-space micromachined optical switches with sub-millisecond switching time for large-scale optical crossconnects," *IEEE Photonics Technology Letters*, April 1998, pp. 525-528.

[11] R. T. Chen, H. Nguyen, and M. C. Wu, "A low voltage micromachined optical switch by stress-induced bending," *12th IEEE International Conference on Micro Electro Mechanical Systems,* Orlando, FL, 1999.

[12] B. Behin, K. Y. Lau, and R. S. Muller, "Magnetically Actuated micromirrors for fiber-optic switching," *Solid-State Sensor and Actuator Workshop*, Hilton Head Island, SC, 1998.

[13] H. Laor, "MEMS mirrors: application in optical cross-connects," *IEEE LEOS Summer Topical Meetings: Optical MEMS*, Monterey, CA, 1998.

[14] D. T. Neilson, V. A. Aksyuk, S. Arney, N. R. Basavanhally, K. S. Bhalla, D. J. Bishop, B. A. Boie, C.A. Bolle, J.V. Gates, A.M. Gottlieb, J.P. Hickey, N.A. Jackman, P.R. Kolodner, S.K. Korotky, B. Mikkelson, F. Pardo, G. Raybon, R. Ruel, R.E. Scotti, T.W. Van Blarcum, L. Zhang, C.R. Giles, "Fully provisioned 112x112 micro-mechanical optical crossconnect with 35.8 Tb/s demonstrated capacity," *Optical Fiber Communication Conference*, Baltimore, March 2000, paper PD-12.

[15] A. Neukermans and R. Ramaswami, "MEMS technology for optical networking applications," *IEEE Communications Magazine*, January 2000, pp. 62-69.