Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

MEMS in the Optical Network

Optical fiber began invading the network in the late 1970s. We all know the fibersphere is building rapidly. I've seen estimates of from fifty to seventy-five thousand miles of fiber going into the ground every day (see, for example, "Switching the Light Fantastic" in *EDN Magazine* 26 October 2000). We're moving from analog switches on all-copper connections to something optical. MEMS will be a key part of the answer. There's a challenge, however, in seeing what the optical network will look like and in seeing how we will get from today's packet-switched rings to the future's fiber-switched grid. I'd like to do that without learning phone acronyms, like RBOC, EAEO, HDBH, CLEC, or ILEC, IXC, NANP, LATA, and POP. Some explanation is necessary, however, to understand where the network is headed and why. For simplicity, I'll divide the network into three pieces: local access, metropolitan (the metropolitan-area network), and long haul (the network core). "Clients" (that's us, for example) connect to the local-access network. The local-access networks are the bridge between analog clients and the all-digital rest of the network. The metropolitan-area network connects among the local-access networks and also connects to the long-haul network. The long-haul network connects metropolitan-area networks.

The network core is changing first, the metropolitan-area networks will change next, and finally, the local-access networks will change. Lower cost and huge increases in demand for bandwidth drive changes in the network core. Today's optical networks take the "business route" through every town on the journey. In the future, the network core will be more like the interstate highway system; local traffic will take the business route and through traffic will have controlled-access and controlled-exit lanes and will bypass towns along the route. Microelectromechanical systems bring economies of power, cost, and batch fabrication to network components. To see where we're going, it helps to look at how we arrived where we are today.

Development of the packet-switched digital network

Wireline voice communication might have begun with two tin cans and a string. This primitive analog communication system was simple and functional, but it didn't scale well since each pair of users needed its own interconnecting string. It also didn't work for long distances.

Phone systems overcame difficulties in scaling by building *switches*: instead of running lines between each pair of stations that might want to communicate, all lines ran to a central switch where any two "subscribers" could be connected. In the labor-intensive early days of phone systems, an operator at the central-office end of the caller's copper wire plugged the caller's line into the desired destination line. A Kansas City undertaker named Almon Strowger ended those days in 1891—when he invented the electromechanical telephone switch. His objective was to displace local operators, whom he suspected of diverting calls to a rival. Strowger's analog, sequential, mechanical circuit switch gave way to the electromechanical crossbar in the 1940s.

Electronics in the form of speakers, microphones, amplifiers, and filters helped overcome distance limitations. An amplifier boosts the energy in a signal, allowing the signal to be propagated to and detected at much greater distances. For analog signals, one doesn't know what the signal is supposed to look like. That is, we don't know which part of the signal is noise and which part is the real signal. So the problem with amplifying an analog signal is that the noise is also amplified and cannot subsequently be separated. Each time an analog signal is amplified, noise accumulates and degrades the signal. Converting the network from analog to digital could solve this problem.

In This Issue:

Conversion from analog to digital began at the core of the telephone network in the mid 1970s. Here's how it works.

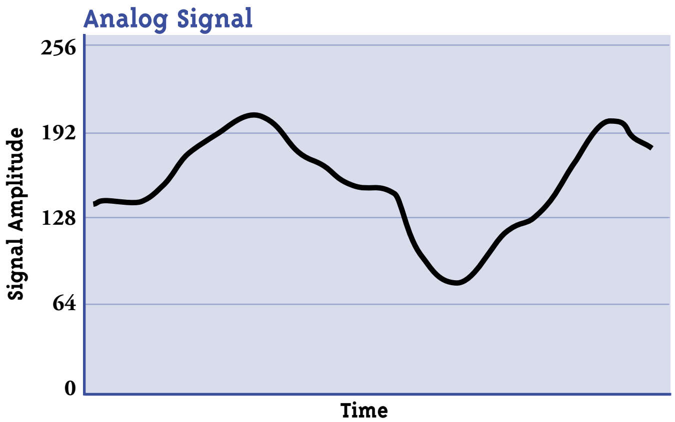

Fig. 1. A one-millisecond sound from Bonnie Raitt's "Baby Come Back."

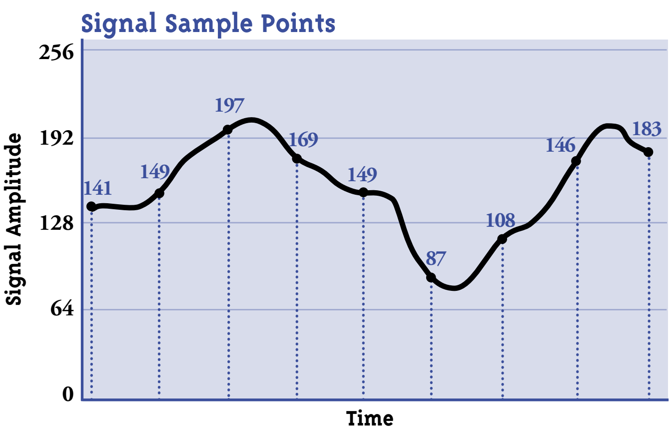

Fig. 2. The amplitude of the analog signal is "sampled" at regular intervals. Each sample, shown in this figure as a decimal value, becomes an eight-bit binary value.

Fig. 1 above is the analog representation of one millisecond of sound from Bonnie Raitt's "Baby Come Back."

To convert "Baby Come Back" to a digital representation, we need to "sample" the signal and save the values. I show such sample points in fig. 2.

If we collect enough samples, it will be possible to approximate the original analog signal. I've illustrated this process in fig. 3. Each sample point is converted into a binary representation of its amplitude (in this case eight bits). We can amplify this string of ones and zeros and can send it to a remote location. We will amplify noise with the ones and zeros, but it will take a lot of noise to flip the value of a bit. At the destination, we can recover the ones and zeros and throw away the noise. Digital signals do not suffer from noise accumulation. We converted the original analog signal to a digital representation (a collection of eight-bit values representing the sample points) so that we could prevent noise accumulation as we amplified the signal through the network, but we gave up something of the original signal in doing so (compare the signals in figs. 1 and 3).

To convert voice to a digital representation, the phone system breaks each second into 8,000 segments and records a volume value between 0 and 255 for each segment (this requires eight bits). The data rate for a voice line, therefore, is 8,000 samples per second at eight bits per sample. This is 64,000 bits per second. A sample rate of 8,000 per second is eight samples per millisecond, which is what is illustrated in figs. 1 through 3.

Ever since it began, the load on the phone system has been increasing. In the days of copper wires and analog electronics, expanding the system meant digging trenches and burying cables. Digital signaling offered an alternative called multiplexing—a technique that allowed conversations to share a single wire. Your plain old telephone service (POTS) line digitizes to a bit rate of 64 kb/s. Once it is converted to a digital signal, it can be multiplexed with other digital signals on a higher-capacity line. The usual method for sharing is time-division multiplexing (TDM). At twice the data

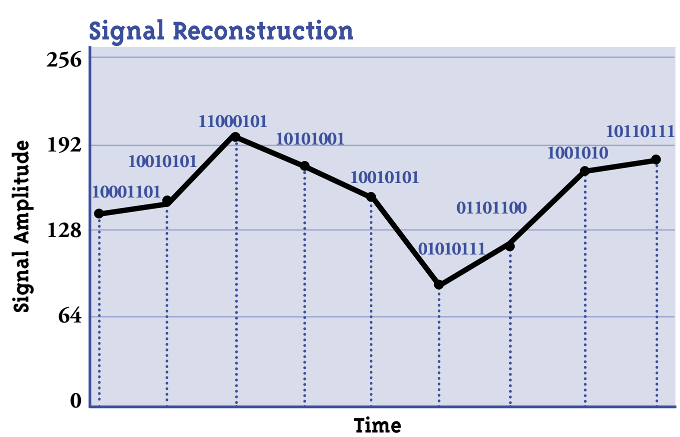

Fig. 3. The transmitter sends the binary ones and zeroes for each sample point to the outbound fiber. The receiver reconstructs the signal from the ones and zeroes of the sample points.

DynamicSilicon

Editors	**Nick Tredennick**
	Brion Shimamoto
Publisher	**Lauryn Franzoni**
Associate Editor	**Keri Geiger**
Designer	**Julie Ward**
President	**Mark T. Ziebarth**
Chairman	**George Gilder**

DynamicSilicon is published monthly by Gilder Publishing, LLC. Editorial and Business address: P.O. Box 660, Monument Mills, Housatonic, MA 01236. Copyright 2001, Gilder Publishing, LLC. Editorial inquiries can be sent to: bozo@gildertech.com Single-issue price: $50. For subscription information, call 800.229.2573, e-mail us at dynamic@gildertech.com, or visit our website at www.dynamicsilicon.com

rate, chunks of two separate signals can be interleaved like the two halves of a zipper. Four times the data rate supports four interleaved signals, and so on. Twenty-four simultaneous conversations can be transmitted over a T1 line with an aggregate data rate of 1.544 Mb/s (twenty-four eight-bit samples plus one framing bit each 125 microseconds)—about the speed of a 1x CD-ROM. Faster lines aggregate more conversations. Table 1 provides telecommunication line speeds and line aggregations.

Table 1. Telecommunications line speeds and line aggregations

Name	Speed	Aggregation	Remarks
DS-0	64 kbps		Plain old telephone service
T1/DS-1	1.54 Mbps	28 POTS	About the transfer rate of a 1x CD-ROM
T3/DS-3	44.74 Mbps	28 DS-1s	High-end office connection
OC-1	51.84 Mbps	28 DS-1s	T3 plus SONET overhead
OC-3	155.52 Mbps	3 OC-3s	
OC-12	622.08 Mbps	4 OC-3s	
OC-48	2,488.32 Mbps	16 OC-3s	"2.5 Gb/s"
OC-192	9,953.28 Mbps	4 OC-48s	"10 Gb/s"
OC-768	39,813.12 Mbps	4 OC-192s	"40 Gb/s"

Moore's law and the Internet

Fig. 4 projects the supply (i.e., link capacity) and demand for bandwidth; it includes Moore's law for comparison. The link capacity projection assumes continuation of a trend from deployment of 2.5 Gb/s links in 1991 through 400 Gb/s (multi-lambda) links in 1998. Demand for bandwidth assumes capacity requirements double every year. These trends are reported in a white paper from Lucent ("Evolution of the Transport Network" by Deirdre Doherty, et al). Most estimates, including those from *IEEE Spectrum*, *IEEE Computer*, and another white paper from Lucent, have bandwidth demand doubling every six months. *EETimes* (27 Sep. 2000, p. 30), quoting another Lucent source, estimated doubling time at nine months. For my purpose, it doesn't matter whether the doubling rate is the likely six months or the conservative year.

Moving from copper to fiber, more options exist for upping the bandwidth. With copper, the options were either to bury more cable or to speed up the data rate. Add lanes or increase the speed limit. Fiber allows multiplexing several *wavelengths* ("lambdas") in the same fiber *and* it boosts the potential data rate for each lambda (since photons are faster than electrons). For increasing link capacity, the costliest proposition is trenching to add new fiber. With fiber, once it's in the ground, we can add lambda lanes and up the speed limit without digging.

Moving bits faster is the traditional way to speed things up, but here we run into Moore's law and its electronic elements. Elements in the transport network follow the link-capacity curve, while switches and routers and the encoders and decoders needed to multiplex information onto higher data rates follow Moore's law in fig. 4. Switches and routers are getting faster at a rate governed by Moore's law, but not fast enough to keep up with demand. Further, mere electronics will *never* be fast enough to exploit the full data-carrying capacity of optical frequencies. The fastest CMOS processors of 2001 operate below 2 GHz. Signals in the range of 1,500 nanometers, popular in optical networks, represent frequencies near 200,000 GHz. That leaves adding wavelengths to the fiber links as the best way to increase link capacity. Adding wavelengths to the fiber is called *wavelength-division multiplexing* (WDM). Each wavelength is a channel and each channel may itself be a time-division multiplexed aggregate of interleaved packets from many virtual connections. In the bandwidth of a fiber between 1,260 and 1,620 nanometers, there's room for more than 500 channels at 100-GHz spacing (and four times that number at 25-GHz spacing). If each channel is modulated at 40 Gb/s, then the data rate carried by an 864-fiber bundle would be 18 Pb/s (Petabits per second, which is 10^{15}). That's enough link capacity to handle today's worldwide data traffic in a single fiber bundle. The range of lambdas that can be transmitted will increase with improvements in transmitters, receivers, and fiber.

When the network was analog, it was circuit switched, meaning that when you called someone, the network made a dedicated physical path, via switches and wires,

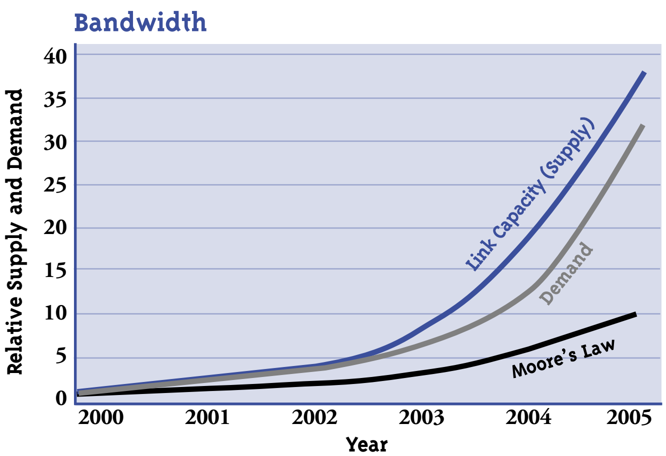

Fig. 4. Projected supply and demand for bandwidth normalized to one in 2000.

between your phone and their phone. When the core of the network became digital, the core became packet switched, meaning that when you called someone the network made a "virtual" circuit, via switches and wires, between your phone and their phone. It's a virtual circuit because the packets of digitized communication between the analog end points travel over shared physical paths. The Internet has grown up as an end-to-end, digital, packet-switched network. Your phone conversations and the Internet's data packets *share* the core of the network.

Rapidly rising numbers of routers and switches in the metropolitan-area network and in the network core support the growth of the Internet, but there's a penalty. The number of routers in the path of a virtual circuit has grown, contributing to delay and to congestion. Today's typical message passes through an average of *seventeen* routers. Picture driving to work, even with light traffic, through seventeen tollbooths. Though the network itself may be fiber, switches and routers convert the signal from photons to electrons to be able to read the packet headers for routing and then reconvert the signal to photons for transmission. If it's a WDM fiber, then the channel for each wavelength must be separated, converted, routed, and reconverted.

Each channel requires a separate detector to convert photons to electrons, a router port for packet analysis, and a laser to reconvert the electrons to photons. A spare backs each of these elements. In addition to these conversions, amplifiers in the network core typically convert the lightwave signal to electrons to clean it up and to boost its strength every 80 to 100 kilometers. These distances increase with improvements in lasers, fiber, amplifiers, and detectors.

Suppose we have a sixteen-channel WDM link between Los Angeles and Boston, with each TDM channel operating at 2.5 Gb/s. Increasing demand suggests it's time to quadruple link capacity. We could do this by increasing the channel speed or by increasing the number of wavelengths (channels) in the fiber. To upgrade channel speed from 2.5 Gb/s to 10 Gb/s, we would replace all of the detectors, routers, and lasers (including the spares) at the link ends and at all intermediate links, plus all of the components in the chain of regeneration amplifiers. Alternatively, we could buy forty-eight new sets of detectors, routers, lasers, regenerators and spares (plus bigger multiplexers and demultiplexers) for the entire link route to upgrade to sixty-four-channel WDM. Speeding up the links from 2.5 Gb/s to 10 Gb/s would be an all-or-nothing proposition. That is, since we have to upgrade the routers and switches that deal with all sixteen lambdas, we might as well upgrade all the detectors and lasers as well. Increasing the number of WDM channels, while still expensive because we have to add equipment at every node in the link, can be incremental (we could upgrade say, four or eight wavelengths at a time). While neither alternative is attractive, they are today's choices in network implementation.

Photonic switches

Let's go to Los Angeles and talk about "grooming." Since there are routers and switches all along the route between Los Angeles and Boston, it doesn't matter which channel or time slot a particular virtual connection's packets travel in because all of the packets can be sorted at each intermediate point.

The first routers to see the packets can begin sorting the traffic so that before it leaves Los Angeles, traffic bound for Boston is bundled into common channels. I'm making this up for illustration, but let's look at an 864-fiber bundle leaving Los Angeles. The route between Los Angeles and Boston might go through Phoenix, Dallas, Nashville, and Philadelphia. LA's routers discover that 400 channels worth of the traffic goes to Dallas (or further), 200 to Nashville (counted in the signals to Dallas), 100 to Philadelphia, and 25 go all the way to Boston. It would make sense to group all of the Boston traffic into channels on common fibers, to group all traffic bound for Philadelphia and beyond on common fibers, and so on. That's grooming. These channels don't have to be converted and regenerated at Phoenix, Dallas, Nashville, and Philadelphia. At Dallas, Nashville, Philadelphia, and other intermediate waypoints, we can route Boston traffic through without looking at it.

The problem with this idea is that the routing information for the packets is in the packets themselves—the routing information is "in-band"—so how can a Dallas router know what to do with the channels unless it breaks them out and decodes them? Since the LA routers have already decoded the routing information, they could build separate communication channels that tell the downstream routers what to do. This is called "out-of-band" routing information. With out-of-band information, the downstream routers can direct fiber traffic without having to decode the traffic itself. This opens the way for *photonic switches*. Photonic switches, called OOO (optical-optical-optical in a reference to input-switch-output) switches, have optical interfaces on the input and output ports and switch the fibers in the optical domain.

Today's optical networks use OEO (optical-electronic-optical) switches. OEO switches have optical interfaces on the input and output ports and electronics in the middle for switching and for signal regeneration. The electronics in the middle limits the speed of the switch and it limits the capacity of the fiber since the fiber cannot carry any more wavelengths than the electronics in the middle (of the OEO switch) can handle. Photonic switches eliminate the electronics bottleneck in the middle of the switch. Companies such as Agilent, Atoga Systems, Avanex, Axsun Technologies, Calient, Cerent, Chorum, LightConnect, Lucent, Lynx, MEMX, Network Photonics, Novation, Onyx Microsystems, Optical Micromachines, Qusion Technologies, Sercalo Microtechnologies, Trellis Photonics, and Xros (Nortel), are building systems or components for the photonic network. Many alternatives are competing for a place in the optical network. Photonic switches will displace OEO switches in the core of the network because the OEO switches scale with Moore's law, which is too slow to keep up with the expansion of network.

Calient (www.calient.net), a pre-IPO startup with a design center in Silicon Valley, builds photonic switches. Calient is beginning beta test of its DiamondWave 256 photonic switch. The switch has two sets of 256-fiber ports and occupies half of a standard seven-foot equipment rack. Inside the switch, there is a switching matrix about the size of a sugar cube. Two silicon chips sit inside the sugar cube, one for each 256-fiber port. Each silicon chip contains an array of 256 MEMS movable, crystalline-silicon mirrors. Each mirror pivots on two axes under electrostatic control. Light from a fiber on one port travels across free space to hit a mirror on the first array, then travels to a mirror on a second array, and then travels to a fiber on the second port. Mirror positioning directs light from any fiber on one port to any fiber on the other port. Calient will soon introduce the DiamondWave 1K, that fits a 1,024x1,024 photonic switch into a single seven-foot equipment rack, and it has plans to build a DiamondWave 4K, with two sets of 4,096 ports.

Lucent shipped its WaveStar LambdaRouter in July 2000. Global Crossing announced live network testing of the system in its networks on 26 September 2000. Lucent's LambdaRouter is a MEMS-based OOO switch employing two silicon-chip arrays of 256 polysilicon mirrors set four inches apart. Fig. 5 shows a needle atop part of Lucent's MEMS mirror array. The light-colored disks are the mirrors. Two concentric rings surround each mirror. Each mirror is hinged vertically to the first ring. The first ring is hinged horizontally to an anchor ring. Electrostatic forces deflect the x-or y-axis of each mirror from zero to three degrees, to direct light from a fiber on one port to any fiber on the other port. Each 256-mirror array occupies less than a square inch of silicon. Lucent is working to build a LambdaRouter with two sets of 1,024 ports.

While it may seem that Calient and others are on similar paths in building photonic switches with MEMS-based mirror arrays, there are significant differences among the approaches. Electrostatic forces deflect Calient's mirrors, while electromagnetic forces deflect the mirrors in some designs. Electrostatic deflection may be more efficient because forces result from trapped charges. Once the mirror is moved to its target location, forces exerted by the trapped charge hold it in place without a continuous supply of power. By contrast, electromagnetic forces require continuously flowing current to maintain deflection angles.

Tiny though they are, Lucent's mirrors are larger than Calient's mirrors, so Lucent's mirrors have more inertia and move slower, but could reflect more energy without melting. Gold plating can increase the energy carrying capacity of a mirror. Lucent's mirror arrays are four inches apart, while Calient's mirror arrays are less than a half-inch apart. This may seem inconsequential, but it isn't. The light-carrying core of a fiber is about nine microns in diameter, a small target. This is roughly like aiming a flashlight at a one-meter target that is 1,500 meters distant versus one that is 12,000 meters distant. The distant target is harder to hit and aiming is more sensitive to shock and vibration.

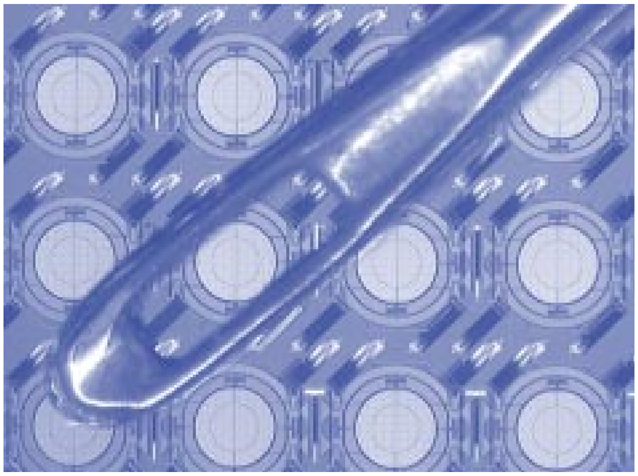

Fig.5. A needle atop polysilicon MEMS mirrors used in Lucent's WaveStar LambdaRouter.

Calient builds its MEMS mirror array with a bulk micromachining process, while Lucent builds its mirror array using surface micromachining. IC manufacture employs two kinds of silicon, bulk silicon and polysilicon. Bulk silicon is the crystal lattice form of the wafer (silicon wafers are sliced from a single gigantic crystal that is grown and purified in an elaborate process). Polysilicon is amorphous; that is, like window glass, it is a collection of particles with no internal structure. Bulk micromachining cuts the mirrors out of the (bulk) crystalline structure, while surface micromachining constructs the mirrors above the wafer's surface through the use of successive steps of deposition and etching. Bulk structures, being crystalline, are rigid, while polysilicon structures tend to be flexible. You have seen the effects of warped mirrors in a fun house; rigid mirrors are better. While increasing the thickness decreases undesired flexibility in the polysilicon mirror, it adds to mass and to inertia.

The likely winner in MEMS-based mirror arrays is small, closely spaced, bulk-silicon mirrors. There are several alternatives to mirror-array switches. All come with tradeoffs. Some, such as thermo-optical and electro-optical switching, use waveguides and diffraction. These switches are sensitive to wavelength and to polarization. Some, such as holographic and liquid crystal switching, use free space. These too are sensitive to wavelength and to polarization, they are expensive, and they may suffer from high insertion loss. Still others use electromechanical positioning or switching. Mirror arrays are insensitive to wavelength, they are insensitive to polarization, they have low insertion loss relative to alternatives, and they are cheap. Alternatives to mirror arrays will find a place in the network where their unique features, such as sensitivity to wavelength or leakage that can be used to monitor the signal, can be used to advantage.

Attractive features of photonic switches include their insensitivity to bit rates and to protocols. The switch doesn't care whether what is being switched is IP, ATM, or SONET and it doesn't matter whether the data rate is 155 Mb/s or 80 Gb/s. Since the switch is insensitive to protocols and to bandwidth, it scales at no cost as the network around it is upgraded. Further, the photonic switch doesn't care how many lambdas are in each switched signal or what wavelengths they are. Wavelengths, which can be added incrementally to any channel without affecting the switch, can be anything in the bandwidth of fiber (1,260 nm to 1,620 nm).

Routers and OEO switches should be pushed to the edges of the network where their electronics are needed to inspect packets. Photonic switches, which incur no electronic processing delay, should displace OEO switches in the core of the network.

The photonic switch doesn't amplify the signal; it only attenuates the signal. The OEO switch amplified the signal when the output laser regenerated it in the process of converting the electrical signal into photons. If the switch doesn't convert the signal to electrons, how will it be amplified? Once again, technology comes to the rescue. The erbium-doped fiber amplifier (EDFA), for example, amplifies all the wavelengths in the fiber by means of laser pumping without having to convert photons to electrons and back to amplify the signals individually. The cost of adding a wavelength to the fiber drops if the EDFAs in the link can accommodate the incremental load.

Optical amplifiers, like the EDFA, amplify signals and noise—the all-optical network behaves more like an analog network than a digital one. Noise will accumulate in the all-optical network just as it did in the analog network. Stronger transmitters, by companies like Princeton Optronics, may reduce the need for EDFA-like optical amplifiers.

MEMS in the optical network

Photonic switches will be based on MEMS, but so will many other components in the optical network. MEMS will invade the optical network because they confer advantages in cost, performance, reliability, and size.

Cost advantages in the MEMS silicon micromachines derive from batch fabrication. "Boats" of twenty to twenty-four wafers go through an assembly-line process. Each wafer may contain hundreds to thousands of silicon micromachines, so the boat contains thousands to tens of thousands. No matter how complex and how intricate the silicon micromachine, its macro-scale equivalent will have to be individually crafted and assembled while the silicon micromachine can be fabricated in tens of thousands.

The size advantages of MEMS come both in component size and in the size of the systems that employ them. Calient's micromirror arrays are a good example of the advantages of being small. Because the micromirrors are small and light, they move quickly, they don't take much energy to move, and wear is reduced. Small, light components are less sensitive to vibration and to shock. In addition, because they are small, tolerances can be closer and aiming is precise.

System size and power consumption are overriding issues for network operators. Remember how the phone company got started by connecting its subscribers to a central office? Those buildings don't move. The network operator's connections to its subscribers are like a giant root system—the network operator can't just pack the furniture and move to a larger building. Meanwhile, the number of subscriber connections is growing and the demands of each subscriber are growing. In addition, the incumbent local exchange carriers must, by law, lease space to competitive local exchange carriers for their equipment. Space at the local carrier incurs a $50,000 one-time charge per central office and about $2,000 per month per hundred square feet. The incumbent doesn't have spare space to lease and the competitive carrier doesn't want to lease more space than is absolutely necessary.

Many of the components in the optical network are candidates for MEMS implementation. Below are some examples of MEMS devices for the optical network.

I'm talking to a friend in Boston. As I speak, the microphone in my telephone handset converts sound pressure waves into an analog electrical signal, amplifies it, and sends it out on twisted-pair copper wires. At the next significant station in its path, the analog signal is converted to its digital packets and interleaved with twenty-three other conversations in a T1 link over copper wire that is bound for a node on the metropolitan-area network. At the metropolitan-area network, the digital packets are groomed and aggregated with other traffic electronically and are then converted to photons for transmission to the network core.

Optical switch. The simplest optical switch either lets the signal pass from one fiber to the next or it deflects the light to a second fiber. Analog control of the switch position results in a variable attenuator. The simple switch is probably the most useful MEMS device in the optical network. Its uses include switching, signal attenuation, power control, add/drop multiplexers, and signal equalization. Switching can rapidly adjust services to meet demand, such as, for example, routing signals around link failures.

Optical add/drop multiplexer (OADM). In addition to the photonic switch, the core nodes require add/drop multiplexers to extract signals for rerouting and to add new signals. The OADM is key to core-network operation. It is how signals enter and leave the core network. There is one OADM at each network node for each fiber strand that can gain or lose lambdas.

OADMs sort inbound and aggregate outbound optical traffic. The OADM needs a way to separate and reroute incoming lambdas. The OADM may use tunable filters to separate inputs and tunable lasers for its outputs. Tuning might employ a diffraction grating or the physical adjustment of the laser's cavity. Dimensions for these adjustments are approximately a quarter of the wavelength. How convenient that wavelengths near 1,500 nanometers amount to physical adjustments of 0.3 to 0.4 microns—within easy reach of standard semiconductor processes, which makes these components shoe-in MEMS applications. The optical network also needs components for variable attenuation, adaptive equalization, and dispersion compensation. All are candidates for MEMS.

Variable attenuator. Short-haul fiber might use fixed attenuators in the line to assure that arriving signals do not swamp receivers that expect low power levels. With WDM signals, a variable attenuator can regulate the power as the number of active signals (lambdas) varies. Variable attenuators adjust the coupling efficiency between two fibers. In one implementation, a flag between the fiber ends moves to block propagation of part of the signal. In a second implementation, a moving mirror adjusts coupling efficiency between the source and destination fibers by aiming the reflection for less than perfect alignment.

Adaptive equalizer. Add/drop multiplexers would mix signals of varying strength on the same fiber. The adaptive equalizer measures the signal strength of each wavelength and adjusts variable attenuators to reduce the variation ("equalize" the signal).

The future optical network

The future optical network switches *paths* (fibers), not data or individual lambdas. Today's networks examine data packets all along the way to determine the path the data is to take. In tomorrow's networks, data packets board a lambda express train all the way to the network's edge.

As the network's traffic load grows, the network will outstrip the ability of electronics to supply the necessary bandwidth. Network links must convert to optical transmission and they must convert to optical switching in a way that can scale independently of Moore's law. OEO switches, that scale with Moore's law because of the electronics in the middle, will not scale rapidly enough. Photonic switches are independent of data rates, protocols, and lambdas in the fiber and will therefore displace electronics-constrained switches in the network's core. OEO switches and routers will remain entrenched at the edges of the network where their electronics are needed to collect, groom, and aggregate outbound traffic and to sort and distribute inbound traffic.

MEMS-based components will support the all-optical network core. Silicon chips with on-board moving mirrors will invade the optical network in photonic switches, simple switches, add/drop multiplexers, variable attenuators, and adaptive equalizers. MEMS-based tunable lasers and tunable filters will also invade the network. These MEMS-based components will be smaller, cheaper, and more efficient than macro-scale equivalents. They are smaller because they are fabricated with semiconductor processes rather than with macro-scale machining. They are cheaper because they are batch fabricated by the tens of thousands on the assembly lines of semiconductor foundries. They are more efficient because they are smaller, have less mass, move smaller distances in tuning, and are on a scale comparable to the wavelengths bring manipulated. There's a compelling case for removing electrical signal conversions in the network core. There's equally compelling logic for the invasion of MEMS throughout the optical network.

Nick Tredennick and Brion Shimamoto
17 April 2001

Dynamic Silicon Companies

The world will split into the tethered fibersphere (computing, access ports, data transport, and storage) and the mobile devices that collect and consume data. Dynamic logic and MEMS will emerge as important application enablers to mobile devices and to devices plugged into the power grid. We add to this list those companies whose products best position them for growth in the environment of our projections. We do not consider the financial position of the company in the market. Since dynamic logic and MEMS are just emerging, several companies on this list may be startups. We will have much to say about these companies in future issues.

Altera and Xilinx (ALTR http://www.altera.com) (XLNX http://www.xilinx.com)
Altera and Xilinx together dominate the programmable logic business, with almost seventy percent of the CMOS PLD market. Both companies are aggressive and competitive. Sixty-six percent of Altera's revenue comes from the rapidly growing communications segment (Telecosm companies) and an additional sixteen percent comes from the electronic data processing (EDP) segment. Altera and Xilinx are positioned to be major suppliers in tethered applications such as the base stations that support mobile devices.

Analog Devices (ADI http://www.analog.com)
Analog Devices is a leader in analog electronics for wireless RF and communication, MEMS for automotive applications (accelerometers, pressure sensors, transducers), and in DSPs.

ARC Cores (ARK (London) http://www.arccores.com)
ARC Cores makes configurable processor cores. Configurable processors allow the application engineer to adapt the processor's instruction set to the requirements of the problem. Conventional microprocessors have fixed instruction sets.

Calient (* http://www.calient.net)
Calient is a pre-IPO startup that builds photonic switches for the all-optical network core. It builds its own MEMS components. Calient has expertise in MEMS components in Ithica, NY through its acquisition of Kionix and through its own experts in Santa Barbara and San Jose, CA.

Cypress (CY http://www.cypress.com)
Cypress Microsystems builds components for dynamic logic applications. Cypress also builds MEMS and is a foundry for MEMS.

QuickSilver Technology, Inc. (* http://www.qstech.com)
QuickSilver has the potential to dominate the world of dynamic logic for mobile devices (untethered). While many companies work on programmable logic and on "reconfigurable computing" for tethered applications, QuickSilver builds adaptive silicon for low power mobile devices.

SiRF (* http://www.SiRF.com)
SiRF builds RF GPS chips for the mobile market. It is a world leader in development of integrated GPS receivers.

Transmeta (TMTA http://www.transmeta.com)
Transmeta makes new generation microprocessors that use closed-loop control to adapt to problem conditions in an x86-compatible environment. This enables Transmeta's microprocessors to save power over conventional microprocessors from AMD and Intel. The base instruction set is not available to the application engineer.

Triscend (* http://www.triscend.com)
Triscend builds microcontrollers with configurable peripheral functions and with configurable inputs and outputs. Triscend helps consolidate the microcontroller market into high-volume, standard chips.

Technology Leadership	Company (Symbol)	Reference Date	Reference Price	3/31/01 Price	52-Week Range	Market Cap.
General Programmable Logic Devices (PLDs)	Altera (ALTR)	12/29/00	26.31	21.44	18.81 - 67.12	10.1B
Dynamic Logic for Mobile Devices	QuickSilver Technology, Inc. (none*)	12/29/00				
MEMS Foundry, Dynamic Logic	Cypress (CY)	12/29/00	19.69	17.73	13.72 - 58.00	2.2B
RF Analog Devices, MEMS, DSPs	Analog Devices (ADI)	12/29/00	51.19	36.24	30.50 - 103.00	14.2B
Configurable Microprocessors	ARC Cores (ARK**)	12/29/00	£3.34	£0.88	£0.76 - 4.57	£499M
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	2/28/01	38.88	35.13	29.80 - 98.31	13.5B
Configurable Microcontrollers (Peripherals)	Triscend (none*)	2/28/01				
Silicon for Wireless RF, GPS	SiRF (none*)	12/29/00				
Microprocessor Instruction Sets	Transmeta (TMTA)	12/29/00	23.50	19.25	12.00 - 50.88	2.3B
Photonic Switches	Calient (none*)	3/31/01				

*QuickSilver, SiRF, and Calient are pre-IPO startup companies.

** ARK is currently traded on the London Stock Exchange

NOTE: This list of Dynamic Silicon companies is not a model portfolio. It is a list of technologies in the Dynamic Silicon paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. The authors and other Gilder Publishing, LLC staff may hold positions in some or all of the companies listed or discussed in the issue.