GILDER TECHNOLOGY REPORT

Published jointly by Gilder Publishing and Forbes Magazine www.gildertech.com October 2000/Vol. V No. 10

Cao's Law

We will get to the companies of the month in due course—opportunities are breaking out all over. But first I would like to introduce the book of the month (and perhaps of the decade; time will tell). It is called *Collective Electrodynamics* and it was written by Carver Mead in his copious free time while launching a revolution in the camera business with the Foveon imager. With a record breaking sixteen million *analog* pixels on a single chip, the Foveon device represents an unprecedented extension to analog of the transistor densities of Moore's Law (predicting a doubling of transistor counts every eighteen months). In analog, where devices do not merely switch on and off but convey information through continuously varying states, chips tend to hold from tens to hundreds of devices rather than the millions that Foveon has contrived. Far superior to all other CMOS (complementary metal oxide semiconductor) imagers, the new chip will be manufactured at 0.18 micron geometries by **National Semiconductor** (NSM), which owns 49 percent of Foveon and which released Richard Merrill, the creator of the fabrication process, to join Foveon. Yet Mead's climactic speech at Telecosm, ending with a prolonged standing ovation, focused less on this amazing new chip and its impact on cameras than on his new book and its promise of a revolution in the physics of the electromagnetic spectrum.

Simon Cao's predictions represent a moment in the history of optics at least as momentous as Moore's law in microchips.

Forty years ago, Mead did the basic research that underlies Moore's Law. Inspired by a 1959 lecture by his Caltech colleague Richard Feynman entitled "There's Plenty of Room at the Bottom," Mead showed that quantum effects previously believed to be a barrier to further miniaturization would actually enable transistors to be made at least a hundred times smaller in linear dimensions than then believed. He concluded that as "the circuits got more complex, they ran faster, and they took less power—WOW! That's a violation of Murphy's law that won't quit."

Also inspired by an idea of Feynman's, Mead's new book may offer a similar promise for the Telecosm and support a new law of technology progress. It recalled to my mind the account of laser inventor Charles H. Townes of a visit to his laboratory in 1956 by the two commanding giants of quantum theory, Nils Bohr and John von Neumann. The two scientists initially told Townes that the laser was impossible because the uncanny coherence of perfectly aligned photons it required violated Heisenberg's uncertainty principle. Mead's new book details how a series of ten experimental developments—from Townes's maser to currents in superconductive rings to quantum wave phenomena at temperatures below one Kelvin—embody a

In This Issue:

coherent state of matter as inconsistent with Heisenberg's uncertainty theory with regard to particles as the laser was. Swept from the stage by this mounting empirical evidence, in Mead's view, are both Heisenberg's theory as a natural law and the photon as a particle.

To Mead, the very hypothesis of a "particle" is an unnecessary legacy of classical physics, when matter was assumed to be made up of solid atoms. Mead's exploration of electromagnetic coherence—perfect continuous alignment of quantum wave functions—gives implicit support to the idea that frequency spacing in fiber optics can be reduced far below the current limits of around 25 GHz. Without obstacles of Heisenberg uncertainty, "There's Plenty of Room at the Bottom" in wave division multiplexing (WDM).

Avanex tests Corvis's limits

This new result might be termed Cao's Law, after Simon Cao of **Avanex** (AVNX). Cao startled the Telecosm conference with a prediction that the number of wavelength bitstreams in a single fiber thread could be increased from **Corvis**'s (CORV) current commercial limit of 160 to the 1000 achieved in Avanex experiments, and then to an eventual level of hundreds of thousands. These predictions represent a moment in the history of optics at least as momentous as Moore's Law in microchips announced in 1965, and in Cao's vision entail the possibility of a "switchless network."

By comparison to the vision of the Avanex sage, even the moving mirrors or tiny bubbles in the best new optical switches seem as advanced and elegant as a nineteenth century railroad, embodied in a track of glass. On a railroad Cao explains, "you have one track," and for the train to go to a different place "the track must move." That is what a railroad switch is, a moving track. That's OK if you have only a few trains.

"But what if you have a lot of trains? Maybe you can build a big switchyard. Maybe you can switch a hundred trains. Maybe a thousand trains. But what if you have a million trains?"

Avanex will rule the world of WDM because it uses the entire panoply of light

Simon is not ready, this day, to announce a million lambdas on a fiber. But more than 1000 WDM lambdas on a fiber adds up to close to a million lambdas on an 864-fiber cable. Moreover, in the lab Avanex has modulated the sidebands of a single lambda channel with 100 different RF subchannels, in effect creating a hundred thousand optical carriers of roughly 120 Mbps on a fiber. Trouble at the train yard. Hundreds of thousands of lambdas in a fiber or tens of millions in a fiber cable cannot be switched like trains.

Even the hope of muxing large "trains" of lambdas together across the switch is dashed by new research from **Corning** (GLW). Eighty percent of data may be "through traffic," not being dropped or added at the node. But the "through" lambdas are scattered through the fibers. Thus, all the WDM signals must be demuxed and switched individually, with through lambdas moved onto "through" fibers. As a result of their research, Corning's own novel switching architecture – based on a sixteen-fiber network with forty lambdas per fiber – calls for 640 cross-connections, one for each lambda. Light all 144 fibers in a Williams long haul cable, however, with current state of the art 160 channels, and that makes 23,040 cross points. At 100,000 channels per fiber that would be…Oh, never mind.

Simon, however, would say that the ultimate objection to optical switches is not the accumulation of problems but the missed opportunity. Why create the flexibility of 100,000 lambdas in a fiber and then not use it at the switch point, which is just where you want it? Why invent the automobile, as it were, and build multilane highways for it, all to rope thousands of cars together and impel them along the highway like a train. "Cars cannot work the way trains work. We cannot have thousands or millions of cars on the highway if the highway must move when a car wants to switch lanes or exit. The car moves, not the highway. On a railroad the track steers the train, but on a highway the car steers itself. *That is why people like cars, they go where you say.* The highway is a network without switches. We have exit ramps or intersections, but no switches."

MFNX lays dark plans for TDM

On the first page of every elementary text on communications networks, the authors explain why networks have switches. In a network with only two nodes, there would be no switches, just a single wire from A to B, my house to your house. You could do a network with four nodes the same way. One wire each from A to B, A to C, A to D, B to C, B to D, etc. Three wires would terminate in every node. Still not so bad. But if there are absolutely no switches, then every phone gets its own wire to every other phone in an ever-expanding tangle. Connect to everyone in a reasonably sized town and the cost of the wire running to your house would be more than the house itself. That's why we have switches: because we have neither the space nor the money for all the wires that would be needed to replace them. Wires have weight, occupy space, and cost money. If wires had no mass and were free we would not need any switches at all. None.

In the Telecosm we are supposed to reverse the law of the Microcosm and spend bandwidth to save switches. Mostly we do, as when we accelerate a creaky packet network by building bigger dumb pipes rather than stuffing the routers with more elaborate Quality of Service orchestration and traffic management. Buying a dark fiber from **Metromedia** (MFNX), rather than T1 or T3 "services" from a telco, we substitute bandwidth for Time Division Multiplex (TDM) processing.

Hollowing out the computer, transforming it into a network appliance, once again we trade bandwidth for processing overhead on a computer's backplane bus.

Nevertheless, it is undeniable that as copper turned into fiber, both packet and circuit networks switched and processed more and more, not less and less. Even at the box level, switching and routing nodes grew into the millions and at the transistor level to the trillions and beyond.

Xros, Calient diffuse paradigm scandal

The source of this paradigm scandal was the use of fiber as merely a fatter TDM pipe. The telcos assumed that *the real physical layer is the fiber rather than the light*. If we focus on the fiber, or regard WDM as a mere multiplier of fiber capacity, we focus on bandwidth to the neglect of connectivity. Connectivity—flexible and virtually infinite links rather than nineteenth century railroad connectivity—comes not from the fiber but the light. It comes from WDM.

As always, to find the paradigm go to infinity and work your way back. Imagine that communication power is limitless and that we can broadcast all the world's information to everyone all the time and that each user has the capacity to sort through the bit streams and find the one meant for him. In that network there is no need to switch, buffer, route or process anything, *except at the terminals.* But organizing that abundant bandwidth into a single shared channel would oblige us to build, at each terminal, the fastest, largest switch the world had ever seen in order to sort through all the world's information to find the bits meant for each user.

If your model blows up at infinity it may be a good idea to tweak it before you get there. As bandwidth grows faster than the number of wires, the switches will become bottlenecks, and our first impulse will be to multiply them to compensate. Only in a network in which the marginal cost and mass of the wires approaches zero will the wires become the prime source of connectivity as well as bandwidth. In that network switches tend to disappear. Lambdas are those wires.

So why aren't switches disappearing?

They are. Optical switches represent the first significant reduction in the amount of switching in the network.

They cut back on switching by reducing the number of ports for any given bandwidth and by handling all bit rates or protocols the same way (thus eliminating separate ports and paths for different bit streams). TDM cannot provide bit rates much above 40 gigs. Even today one port on a **Xros** or **Calient** switch can handle at least six times that, assuming a hundred 2.5 gig lambdas, and is ultimately limited only by the melting point of the mirror.

Switchless connectivity, however, implies virtually infinite, nearly massless wiring, at a nominal cost per wire. That is, it implies dividing the mass and cost of a fiber over such a large number of lambda channels as to make the marginal lambda virtually free.

As Simon says, even today railroads are great for some purposes. If you are moving a lot of freight between just two points, there is no need to steer. Tracks bear a lot of weight, and a few switches—which may spend days or years in one position—provide all the flexibility and connectivity you need. An undersea network link is like a freight train. But if you have a thousand lanes or a million lanes, then you want to steer. And you do not want to be steered by a central network control system. You want to steer yourself. You need what Cao calls a "smart photon."

In invoking the smart photon, Simon is not indulging the vain hope of an optical computer, a box in which photons are tortured until they behave like electrons. "No. You steer by frequency."

Imagine hundreds of thousands of frequency selective optical pathways in every fiber. The paths are defined by WDM frequency channels. "On Lambda One, you always go from A to B. On Lambda Two you always go from A to C. So imagine that I am at A. I have a tunable laser that I can toggle from Lambda One to Lambda Two. If I want to go to B, I toggle One. If I want to go to C, I toggle Two. I never need to ask the network to switch."

Optical switches represent the first significant reduction in the amount of switching in the network

Such a network has already been designed and built on a small scale in Norway by **Telenor**, using tunable lasers from **Altitun** (ADCT), **Marconi**, and **NTT**. Such a limited mesh would incur traffic jams. The path from A to B may be blocked. But that's where the hundred thousand lanes come in. Avanex alone can enable thousands on thousands of lanes of long distance light.

Fiber-borne light can enter the mind only in the form of metaphors. A "channel" implies a chute or a tunnel or a pipeline. To get more of them we "space" them "closer together." Of course the channel is not really a container and frequencies aren't really spatial, so the channels can't really be close together or far apart. They flow together. Nevertheless, the light waves forming the channel can be thought of as having a "shape." Good shape is the key to good WDM.

Roughly speaking, we call a channel by its center frequency, the frequency equidistant on either "side" from the center frequencies of neighboring channels. Surrounding this center frequency is a band of frequencies that are just a bit off center. How many? Well, leading WDM systems today operate at 100 gigahertz spacing between channels. Which means between one center frequency and the next there are at least another 100 billion frequencies, just in case you were afraid we'd run out. The frequency bands close to the center frequency are important because they give

NET ECON 101

Inadequate connection speeds (Chart 1) are not the sole source of network latencies. Close to seventy-one percent of all data packets requested over the Internet either originate in or are routed through the United States (Chart 2), an imbalance that weighs heavily in the World Wide Wait.

Rapidly improving last mile links and effective storewidth strategies such as Exodus IDCs and Mirror Image's global CAP network will ease both these challenges. Nevertheless, as long as the Net remains a shared, statistically multiplexed medium, it will be prone to the most irritating sort of delay: unpredictable and hugely variable increases in latency, traffic jams from too many users tapping into the same pipe.

The lambda network and a return to circuit switching, with its guaranteed bandwidth, are the ultimate answer. Until then the best way to avoid traffic jams is to continue pumping up bandwidth to reduce the utilization levels of the lines (Chart 3). The only thing better than big, dumb pipes is big, dumb almost empty pipes. Low line utilization means more headroom for sporadic bursts in data traffic and fewer traffic jams.

As Gordon Stitt, president of Extreme Networks, pointed out at Telecosm, the economics of the Internet are determined by speed. Network latency is inversely proportional to end user productivity. A miniscule reduction in network response time from two seconds to one second could effectively double end user productivity in a highly transactive B-2-B Web business. Zona Research estimates that in 1999, when e-commerce was just beginning to take off, e-businesses were losing some $360 million per month in online sales because of disappointing download speeds (Chart 4).

Faced with those numbers what IT manager would not jump at Cogent's 100 Mbps line for $1,000 per month, 70 times the speed of a T-1 for the same price. Bandwidth glut? I don't think so.

– Mary Collins

Sources: CAIDA, ISP World, IBM, Zona Research

Network speeds have suffered due to insufficient Internet connection speeds and...

Estimated connection speeds of U.S. Internet users

Chart 1

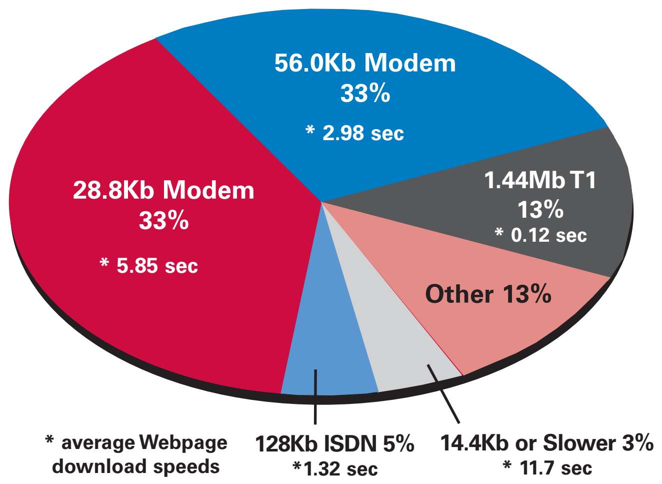

...an imbalance in the amount of data packets routed through the U.S.

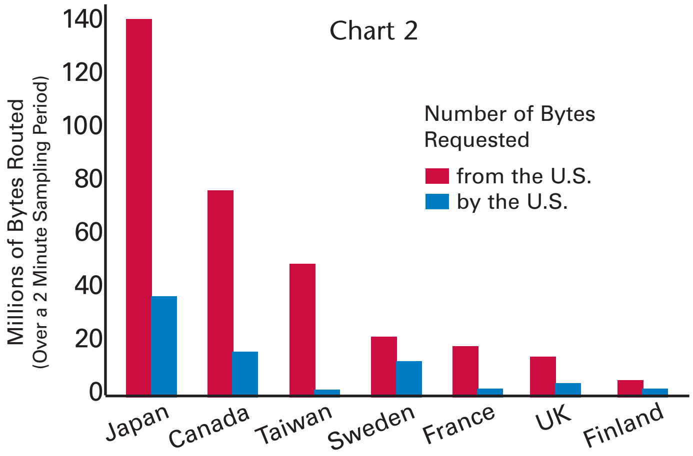

A surge in optical bandwidth will lower percent line utilizations...

Traffic flow from U.S. to Swiss Academic Research Network

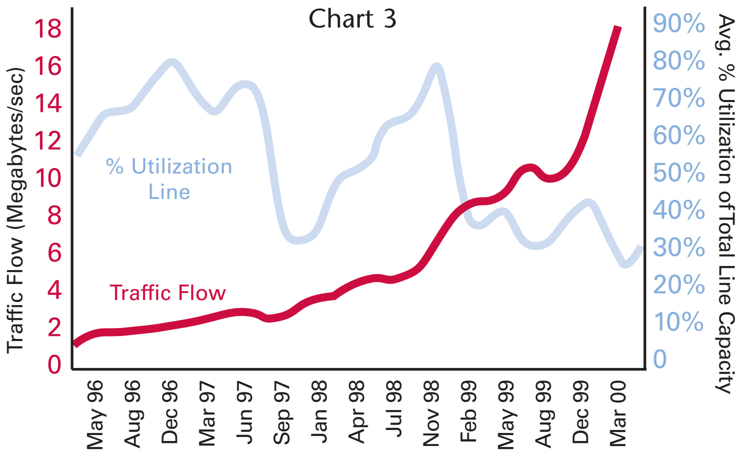

...improve Network latency, and boost web economics.

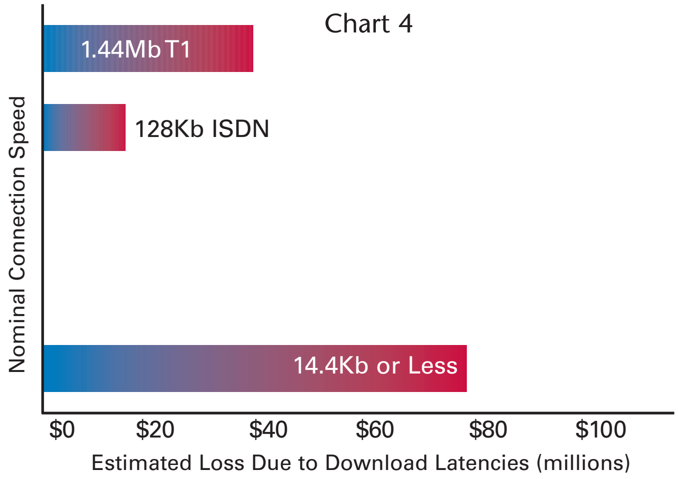

our lambda channel oomph! You can't have a signal without signal power, and lots of the power is distributed just off either side of the center frequency. Altogether the center frequency and its relatively near neighbors are called the "passband."

Avanex banishes bell bottoms

Across the passband there will be an unfortunate tendency for the amplitudes to slope down and away in a Gaussian bell curve, spreading our signal power into a broad-based bell-bottomed channel. A WDM system that spreads significant signal power out to its bell bottoms, say 50 GHz off center, will just barely accommodate 100 GHz spacing, about one hundred channels in a standard fiber. So no bell bottoms. We want channels that look like shoeboxes stood on end, which allow you to approach the maximum bandwidth per channel because more signal power is focused "on target," in the identifiable frequencies of our WDM network.

How do we shear and shape photonic channels? In addition to wavelength and frequency, the language of light also offers polarity and phase. Waves that are aligned in phase add, or interfere constructively, strengthening the light stream; waves that are out of phase cancel, or interfere destructively. Waves of different polarity can take different paths. Avanex will rule the world of WDM because it uses the entire panoply of light.

Before Avanex the three contenders for dominance of the WDM market were thin-film filters (TFF), arrayed waveguide gratings (AWG), and fiber Bragg gratings (FBG). In these technologies, resolution of frequency channels depends on precise *spatial* alignment, difficult to achieve and sensitive to heat and other environmental factors. After enhancing commercial WDM channel count more than forty fold and bandwidth several thousand fold in five years, all three techniques now face show-stopping manufacturing challenges that Avanex's PowerMux deftly avoids.

The PowerMux is based on a sophisticated adaptation of the Fabry-Perot interferometer. A classical device that was tested in some early WDM systems, Fabry-Perots were deemed too slow for packet switching and gave way to thin-film filters which are based on similar principles. The core of a Fabry-Perot is an etalon—two highly reflective mirrors facing each other across a cavity. The input signal enters the etalon cavity through the back of one of the mirror surfaces. It then traverses the width of the cavity to the other mirror where a small portion, usually less than 5 percent, of the beam passes through and the rest reflects back to the rear mirror where part is again reflected and part passed through. As these reflections continue back and forth between the mirrors, the light fluxes leaving the filter cavity emerge at different points in their wave cycles, adding in phase for those frequencies which are related to multiples of the one-way propagation delay across the cavity, creating coherent passbands.

Thus, unlike other commercial WDM filters—which split light along different pathways that must be spatially aligned—frequencies in the Fabry-Perot-based PowerMux are self-aligning because the interference patterns are established by the different wavelengths themselves as they repeatedly traverse the same path. "Align one frequency," as Cao explains, "and the other equally-spaced frequency channels automatically line up." The light does the work.

In the Microcosm, self-aligned silicon gates, invented at Fairchild and perfected at **Intel** (INTC), were critical for the development of integrated circuits. Self-alignment will be similarly crucial to the triumph of WDM. It enables high channel counts because it cancels out the manufacturing showstopper—spatial alignment—that would practically limit the multiple path devices to a few hundred lambdas anywhere outside the lab.

Avanex's PowerMux deftly avoids the manufacturing challenges of other muxing techniques

Transforming the linear Fabry-Perot device into an asymmetric non-linear function, Cao can build exceptionally regular rectangular channels with minimal spacing; or in the different design of the PowerShaper can pre-distort the wave to adapt to a particular fiber dispersion profile. Insensitivity to both the spatial and wavelength domains make the PowerMux and the related PowerExchanger (add-drop) scalable to channel counts limited in number chiefly by the resolution of lasers and other components.

Simply by altering the gap between its etalon mirrors a Fabry-Perot could be tuned to yield lambdas across the entire spectrum of optical fiber. Even a tunable AWG would have a far more restricted range, limited by those multiple aligned paths which pre-determine the frequency intervals, and the challenge of heat tuning a device with so many coordinated elements.

AWGs and other multiple path devices are good for pitching high and hard, shooting a few lambdas at high bit rates and high power. But in the Telecosm as in the Microcosm, "low and slow," and "wide and weak," win the day. Operating across the entire available spectrum the PowerMux will multiply channels of nearly any spacing and size, in the future slicing lambdas down to megabit proportions or even less if we like.

With nonlinearities rising by roughly the square of the power, the tighter the channel the lower the power, the better the optics work. Eat your heart out, Murphy. Ultimately, networks bearing a hundred thousand lambdas per fiber will have lower bit error rates and require less complicated and costly contrivances to maintain signal quality than sixteen or thirty-two lambda systems require today. Most of the problems that plague fiber transmission result from the need to maintain coherent passbands tens or hundreds of billions of Hertz wide. The narrower the band, the less power in the drifting bell bottoms, the fewer the distortions entailed in differential wavelength response to the medium. The less the space, the more the room. It would all sound familiar to Carver Mead or to Richard Feynman.

Corning's coup

Moore's Law, as enabled by Mead's insight, functioned for two decades before its implications were accepted even by the computer industry. In fiber optics, outside Avanex, even the titans of the Telecosm still succumb to the temptations of high and hard, pushing bit rate over channel count. Visiting Corning last week–swiftly becoming along with **JDS Uniphase** (JDSU) our favorite broad-based components company–we listened to Wendell Weeks, Corning's EVP of Optical Communications, enthuse about Corning's conquest of 10 Gbps and 40 Gbps transmission. The maximum allowable dispersion at 40 Gbps, he gleefully tells us, is a mere four one thousandths of what can be tolerated at 2.5 Gbps. (Translation: signal quality control has to get 250 times better to sustain a 15 times increase in speed.)

Weeks loves this problem, because he thinks Corning by virtue of both its prodigious research effort (at 8 percent of sales it leads the industry) and its position as both the world's leading fiber maker and now an industry leader on the components side, is uniquely positioned to solve it for everyone. Crucial to Corning's strategy is its hybrid EDFA/Raman amplifier.

As with most new components for high data-rate systems, Raman amps react differently to different fibers. So Corning is developing a smart card that will detect the connecting fiber and adjust the gain appropriately. A new generation of variable gain EDFAs will push customization one step further adjusting for different spacing in different networks. The tweaking also works from the fiber side with new hybrid fibers bearing dispersion profiles finely tuned to the sensitivities of high bit rate networks. These heroics of integration, perhaps only possible at Corning, would create high barriers to entry for rivals.

Along with JDSU, Corning is swiftly becoming our favorite broad-based components company

The danger for Corning is that they may actually be smart enough to do it. Companies that excel at engineering their way through problems are often tempted to turn problems into business plans, engineering a "gotcha" strategy that holds customers hostage to a self-fulfilling prophecy of problematically optimized proprietary systems. The temptation of speed freaks everywhere, the "gotcha" strategy has killed any number of companies afflicted by engineering narcissism, including Cray and nearly **IBM** (IBM).

Still, don't bet against Corning. If they can escape the TDM temptation, they can seize the day in WDM. Moving optical component manufacturing out of the craft-guild era and into the industrial age represents one of the Telecosm's greatest challenges. Corning commands as many patents in process technology as in the famous materials. Corning's joint venture with Samsung to automate the manufacture of thin-film filters has already improved production yields by some 25 percent and reduced cycle times from two weeks to two days. Even in a PowerMux world thin-film filters will likely remain a crucial WDM technology for coarser separations.

Already the world's best fiber maker, Corning is successfully playing catch-up in the one area Lucent was ahead. Like Lucent with its AllWave, Corning has eliminated the water spike in its MetroCor fiber, opening up the 1400 nm region, expanding the passband by 40 percent. Holding back AllWave, EDFAs do not operate at 1400 nm, limiting transmission distance to about 40 km. Corning is at work on an amplifier (probably using thulium) that could open up the 1400 nm region for long-distance transmission. Along with JDS Uniphase, Corning can lead the way to cheap components for a ubiquitous fibersphere. But it must grasp the imperative of low and slow—low powered signals in such profusion that fast packet switching becomes entirely unnecessary in the core of the network.

Cypress sprouts

In the Microcosm, one of the pioneering advocates of "low and slow" was T. J. Rodgers who formed **Cypress Semiconductor** (CY) in the early 1980s to pursue CMOS at a time when it was perhaps one thousand times slower than its mainstream rivals. An expert on semiconductor processes who invented a unique trench technology while still at Stanford, Rodgers continued to innovate through the 1990s. But Cypress missed most of the key communications chip markets and fell short of the Telecosm.

Within the last year, however, the company has made a dramatic turnaround. Today two thirds of its sales are to communication systems companies, from Nortel to **Motorola** (MOT). Early this month, along with record quarterly revenues of $356.2 million, Cypress announced a reorganization into four perfectly telecosmic divisions. It has adapted its industry-leading portfolio of fast static random access memories (SRAM) to an array of specialized communications slots. Coupled with digital signal processors in mobile phones and other wireless appliances, its MoBL SRAMs are optimized to conserve power and thus prolong battery life. Its NoBL synchronous SRAMs provide the higher bandwidth needed in new data-intensive 3G base stations. It has developed programmable spread spectrum clock chips that dramatically reduce electromagnetic noise and are vital to networking and graphics applications. This summer, Cypress capped off its move to communications by purchasing Silicon Light Machines, a maker of optical MEMS devices usable in a range of WDM applications. Rodgers had served on SLM's board, and Cypress had provided a foundry for its chips and had invested in the company. Since SLM's device can be produced in Cypress's CMOS fabs, these "ribbon-based MEMS" enable the integration of logic with mechanical structures such as gratings, tunable filters, and optical add-drop multiplexers. We will have more to say in future issues. But Cypress's turnaround is not a secret and we hasten now to add them to the list.

EXDS and GBLX consumate courtship

The optical bandwidth tsunami is coming. At Telecosm, **Exodus**'s (EXDS) Niel Robertson reminded us that there is no safer place to be when the wave reaches shore than inside an Exodus Internet data center (IDC). These Telecosm Arks are equipped with dedicated customer bandwidth, multiple backup generators, seismically braced racks, and security worthy of CEO Ellen Hancock's IBM roots. Now in an elaborate deal with **Global Crossing** (GBLX), Exodus is buying GlobalCenter and acquiring discounted access to the GBLX network.

With fifty-one data centers astride major network access points and an industry leading 16 percent global market share in web hosting, Exodus will also become the most connected player on the storewidth seas and will move toward dominance in Asia [where **Asia Global Crossing** (AGCX) already commands 60 percent of all Japanese international IP traffic].

Storewidth Value Matrix

Chart 5

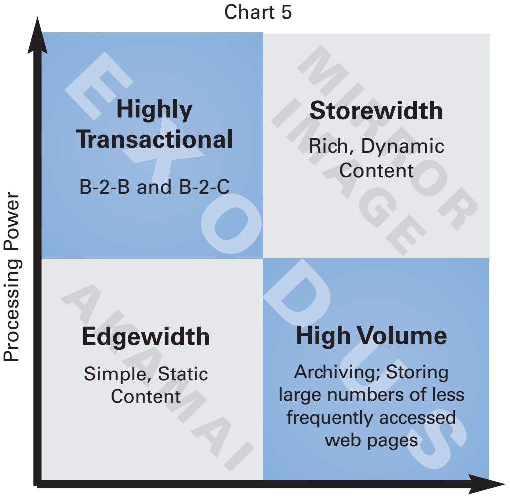

Source: Exodus, "The Storewidth Battlefront", Telecosm

Along with partner **Mirror Image** (XLA), Exodus has discovered the golden ratio of storewidth—the optimal balance of the four key storewidth variables, cache efficiency, management efficiency, network latency, and computer power. Cache efficiency rises with the number of users linked to a cache. Multiplying cache sites to get closer to customers results in fewer user connections. Cache efficiency dwindles. Dispersion of caches also exacerbates problems of management complexity and cache coherency, which is consistency among different versions of the same changing content.

For a dynamic ever-changing Web, cache coherency is the crucial challenge. Failure to achieve it doomed the massively parallel super computer, which on a smaller scale faced all the problems of storewidth currently confronting the Net. Like fashionable net accelerators from **Akamai** (AKAM), **Axient**, and **Adero** today, once fashionable machines from Cray, IBM, and Thinking Machines fell for the temptation of distributing memory close to each processor in order to overcome light-speed delays from central storage. The result was tremendously fast processing of a few specialized programs that mapped well to the massively parallel topography, but crippling incompatibilities among different versions of memory contents for all other applications.

In response to a combinatorial explosion of content conflicts, computer architects contrived ever more Byzantine mazes of buffers and caches. In the end, distributed memory failed because the costs in cache incoherence and complexity dominated the gains in processing speed. The "edgewidth" strategy (Chart 5) gave way to a storewidth strategy.

On the Internet, where the same content may be addressed and altered around the globe, a similar outcome is likely. As Internet contents move up in volume and in processing demands for sorting, searching, and modifying (up and to the right on the chart), storewidth companies such as Exodus and Mirror Image will prevail. Ironically, the Exodus solution also will end up suffering negligibly if at all in latency or delay.

The prime measure of latency is the number of hops from the data source to the customer. Edgewidth companies minimize hop counts by cozying up to the end users but end up multiplying cache sites into the thousands and reducing the amount of cacheable material. The Mirror Image/Exodus system, using a few dozen content access points, maximizes cacheablity at a cost of being, on average, just one hop further from the end user.

Exodus's services are crucial to ISPs, storage service providers (SSPs) like **Storage Networks** (STOR), and content providers, who live inside of Exodus IDC cages under the protection of Exodus service level agreements. In fact, Storage Networks is one of Gilder Publishing's new neighbors at the Exodus IDC in Boston, and Tom Casey, Leo Hindery's successor as CEO of Global Crossing, sits on the Storage Networks board of directors. With these estimable allies, Ellen Hancock is creating a storewidth powerhouse perfectly situated to grow with the Telecosm.

George Gilder & Richard Vigilante
with Charles Burger
October 18, 2000

TELECOSM TECHNOLOGIES

ASCENDANT TECHNOLOGY	COMPANY (SYMBOL)	REFERENCE DATE / PRICE		SEP '00: MONTH END	52 WEEK RANGE	MARKET CAP
WINGS OF LIGHT						
Wireless, Fiber Optic Telecom Chips, Equipment, Systems	Lucent (LU)	11/7/96	11 25/32	30 7/8	28 1/16 - 84 3/16	103.1B
Wave Division Multiplexing (WDM) Systems, Components	Ciena (CIEN)	10/9/98	4 9/32 †	122 13/16	14 11/16 - 136 1/4	34.9B
Wireless, Fiber Optic, Cable Equipment, Systems	Nortel (NT)	11/3/97	11 1/2	60 1/4	23 3/16 - 89	179.3B
Optical Fiber, Photonic Components	Corning (GLW)	5/1/98	40 15/16	300	63 3/4 - 340	88.2B
Wave Division Multiplexing (WDM) Components	JDS Uniphase (JDSU)	6/27/97	3 5/8	94 11/16	27 1/4 - 153 3/8	74.1B
Adaptive Photonic Processors	Avanex (AVNX)	3/31/00	151 3/4	107 11/16	47 3/8 - 273 1/2	6.9B
All-Optical Cross-Connects, Test Equipment	Agilent (A)	4/28/00	88 5/8	48 15/16	38 3/16 - 162	22.2B
THE LONGEST MILE						
Cable Modem Chipsets, Broadband ICs	Broadcom (BRCM)	4/17/98	6*	243 3/4	53 - 274 3/4	54.1B
S-CDMA Cable Modems	Terayon (TERN)	12/3/98	15 13/16	33 15/16	18 7/8 - 142 5/8	2.1B
Linear Power Amplifiers, Broadband Modems	Conexant (CNXT)	3/31/99	13 27/32	41 7/8	26 1/2 - 132 1/2	9.5B
THE TETHERLESS TELECOSM						
Satellite Technology	Loral (LOR)	7/30/99	18 7/8	6 1/8	5 - 25 3/4	1.8B
Low Earth Orbit Satellite (LEOS) Wireless Transmission	Globalstar (GSTRF)	8/29/96	11 7/8	8 5/8	5 13/16 - 53 3/4	835.9M
Code Division Multiple Access (CDMA) Chips, Phones	Qualcomm (QCOM)	7/19/96	4 3/4	71 1/4	45 5/16 - 200	53.1B
Nationwide CDMA Wireless Network	Sprint (PCS)	12/3/98	7 3/16 *	35 1/8	27 13/16 - 66 15/16	32.6B
CDMA Handsets and Broadband Innovation	Motorola (MOT)	2/29/00	56 53/64	28 3/4	27 1/4 - 61 1/2	62.7B
Wireless System Construction and Management	Wireless Facilities (WFII)	7/31/00	63 5/8	57 11/16	30 5/8 - 163 1/2	2.4B
THE GLOBAL NETWORK						
Broadband Fiber Network	Level 3 (LVLT)	4/3/98	31 1/4	77 1/8	49 7/8 - 132 1/4	28.3B
Broadband Fiber Network	Metromedia (MFNX)	9/30/99	12 1/4	24 5/16	11 7/8 - 51 7/8	13.4B
Submarine Fiber Optic Network	Global Crossing (GBLX)	10/30/98	14 13/16	31	23 3/8 - 61 13/16	27.3B
Broadband Fiber Network	NEON (NOPT)	6/30/99	15 1/16	34 7/8	27 7/8 - 159	580.9M
Telecommunications Networks, Internet Access	WorldCom (WCOM)	8/29/97	19 61/64	30 3/8	25 1/4 - 61 5/16	87.3B
CACHE AND CARRY						
Directory, Network Storage	Novell (NOVL)	11/30/99	19 1/2	9 15/16	7 7/8 - 44 9/16	3.3B
Java Programming Language, Internet Servers	Sun Microsystems (SUNW)	8/13/96	13 3/4	116 3/4	43 3/4 - 129 5/16	185.6B
Network Storage and Caching Solutions	Mirror Image (XLA)	1/31/00	29	19	2 11/16 - 112 1/2	2.0B
Disruptive Storewidth Appliances	Procom (PRCM)	5/31/00	25	29 9/16	6 3/4 - 89 3/4	337.5M
Remote Storewidth Services	Storage Networks (STOR)	5/31/00	27*	102 3/16	79 1/8 - 154 1/4	9.3B
Complex Hosting and Storewidth Solutions	Exodus (EXDS)	9/29/00	49 3/8	49 3/8	15 3/32 - 89 13/16	20.6B
THE MICROCOSM						
Analog, Digital, and Mixed Signal Processors	Analog Devices (ADI)	7/31/97	11 3/16	82 9/16	23 5/16 - 103	29.5B
Silicon Germanium (SiGe) Based Photonic Devices	Applied Micro Circuits (AMCC)	7/31/98	5 43/64	207 1/16	27 3/16 - 215 1/4	26.0B
Programming Logic, SiGe, Single-Chip Systems	Atmel (ATML)	4/3/98	4 27/64	15 3/16	7 1/2 - 30 11/16	7.1B
Digital Video Codes	C-Cube (CUBE)	4/25/97	23	20 1/2	14 1/4 - 106 1/4	1.0B
Single-Chip ASIC Systems, CDMA Chip Sets	LSI Logic (LSI)	7/31/97	15 3/4	29 1/4	21 9/16 - 90 3/8	9.1B
Single-Chip Systems, Silicon Germanium (SiGe) Chips	National Semiconductor (NSM)	7/31/97	31 1/2	40 3/4	23 1/2 - 85 15/16	7.3B
Analog, Digital, and Mixed Signal Processors, Micromirrors	Texas Instruments (TXN)	11/7/96	5 15/16	47 3/8	37 7/8 - 99 3/4	77.7B
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	10/25/96	8 7/32	85 5/8	31 7/8 - 98 5/16	28.1B
Seven Layer Network Processors	EZchip (LNOP)	8/31/00	16 3/4	35	3 3/8 - 43 3/4	225.8M
Network Chips and Lightwave MEMS	Cypress Semiconductor (CY)	9/29/00	41 9/16	41 9/16	21 5/16 - 58	5.0B

ADDED TO THE TABLE: CYPRESS and EXODUS

† SPLIT ADJUSTED THIS ISSUE

* INITIAL PUBLIC OFFERING

NOTE: The Telecosm Table is not a model portfolio. It is a list of technologies in the Gilder 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. Mr. Gilder and other GTR staff may hold positions in some or all of the stocks listed.

Gilder Technology Report ***Published by Gilder Publishing and Forbes Inc.***
Monument Mills • PO Box 660 • Housatonic, MA 01236
Tel: (800)292-4380 • Fax: (413)274-3030 • Email: info@gildertech.com

EDITOR: George Gilder
PUBLISHER: Richard Vigilante
TECHNOLOGY ANALYSTS: Charles Burger, Mary Collins, Bret Swanson
MANAGING EDITOR: Debi Kennedy
DESIGNER: Julie Ward

SUBSCRIPTION DIRECTOR: Rosaline Fernandes
GENERAL MANAGER, GILDERTECH.COM: David S. Dortman
PRESIDENT: Mark T. Ziebarth

FOR SUBSCRIPTION INFORMATION TELEPHONE TOLL FREE:
(800) 292-4380
Website:
www.gildertech.com