GILDER TECHNOLOGY REPORT
FEBRUARY 1997 Volume II Number 2
Published Jointly by GILDER TECHNOLOGY GROUP and FORBES MAGAZINE

THE DARK FIBER PARADIGM

In 1996, the new fiber paradigm emerged in full force. Parallel communications in all optical networks became the dominant source of new bandwidth in telecom.

Imagine that in 1975 you knew that Moore's Law–the Intel Chairman's projection of the doubling of the number of transistors on a microchip every 18 months–would hold for the rest of your lifetime. What if you knew that these transistors would run cooler, faster, better and cheaper as they got smaller and were crammed more closely together? Suppose you knew the Law of the Microcosm: that the cost effectiveness of any number "n" transistors on a single silicon sliver would rise by the square of the increase in "n."

As an investor, knowing this Moore's Law trajectory, you would have been able to predict and exploit a long series of developments: the emergence of the PC, its dominance over all other computer form factors, the success of companies making chips, disk drives, peripherals, and software for this machine. With a slight effort of intellect, you could have extended the insight and prophesied the digitization of watches, records (CDs), cellular phones, cameras, TVs, Broadcast Satellites, and other devices that can use miniaturized computer power.

If you did not know precisely when each of these benisons would flourish, you would have known that each one was essentially inevitable. To calculate approximate dates, you had only to guess the product's optimal price of popularization and then match its need for MIPS (millions of instructions per second of computer power) with the cost of those MIPS as defined by Moore's Law.

Merely by using this technique of Moore's Law matching–and holding to it with unshakable conviction for nearly 20 years–I became known as a "futurist." Today I await the death of television, telephony, VCRs, and analog cameras with utter confidence as Moore's law unfolds. You can tell me about the 98 percent penetration of TVs in American homes, the continuing popularity of couch potato entertainments, the effectiveness of broadcast advertising, and the profound and unbridgeable chasm between the office appliance and the living room tube. But I will pay no attention. Just you wait–Jack Welch, Ted Turner and Rupert Murdoch, John Malone and David Jennings–the TV will die and you may be too late for the net.

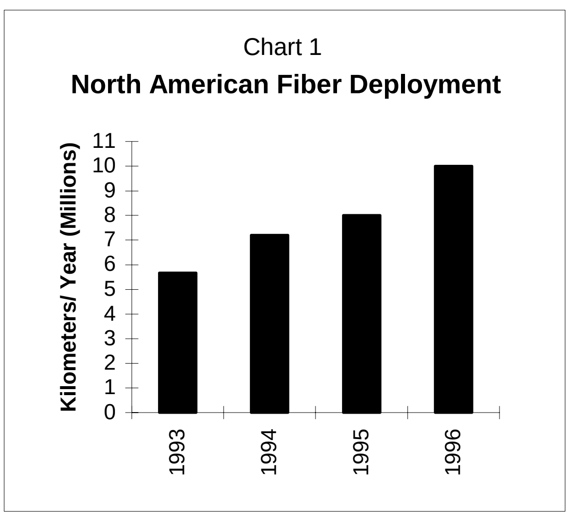

It is now 1997, and a stream of dramatic events certifies that another law, as powerful and fateful and inexorable as Moore's, is gaining a similar sway over the future of technology. It is what I have termed the Law of the Telecosm.

Its physical base lies in the same quantum realm of eigenstates and band gaps that governs the performance of transistors and also make photons leap and lase. But the telecosm reaches beyond components to systems, combining the science of the electromagnetic spectrum with Claude Shannon's information theory. In essence, as frequencies rise and wavelengths drop, digital performance improves exponentially. Bandwidth rises, power usage sinks, antenna size shrinks, interference collapses, and error rates plummet.

The law of the telecosm ordains that the total bandwidth of communications systems will triple every year for the next 25 years. As communicators move up spectrum, they can use bandwidth as a substitute for power, memory, and switching. This results in far cheaper and more efficient systems.

The law of the telecosm ordains that the total bandwidth of communications systems will triple every year for the next 25 years.

Like Moore's Law, the law of the telecosm will reshape the entire world of information technology. It defines the direction of technological advance, the vectors of growth, the sweet spots for finance. Like Moore's Law, it feeds on an array of related advances. Most powerful is Metcalfe's Law. Propounded by Robert Metcalfe, the inventor of the ethernet and a pioneer of the Internet, Metcalfe's Law ordains that the cost effectiveness of networks grows in exponential proportion to the number and power of terminals compatibly attached to them: Take any number "n" computers and link them in a network and you get "n" squared performance and value. Shrewdly adding in the falling price of network adapters and other gear as their volume increases, Metcalfe effectively included Moore's Law in his projection. If you want to understand the runaway growth of the Internet, consider that it feeds on these two exponentials.

The telecosm thrives on low powered bandwidth in the higher reaches of the electromagnetic spectrum. Using infrared waves at 800, 1310, and 1550 nanometers, the exemplary up-spectrum technology is fiber optics. But experts erect a spurious conceptual wall between wireless and wired technologies, which are converging rapidly. For example, **Canon** (CANNY) this year introduced the Canobeam optical network, which takes the same infrared carriers used in fiber and deploys them in a wireless campus area network over several kilometers at fiber like ATM speeds of 155 megabits a second. Meanwhile, in the laboratory, fiber speeds are moving into the terabits (trillions of bits) a second.

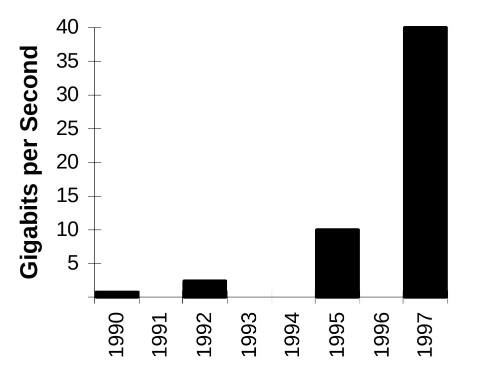

Chart 2

Fiber Speeds in Commercial Use

For more than a decade, American companies have been laying fiber at a pace of some 4000 miles a day, for a total of more than 25 million strand miles. Five years ago, the top 10 percent of US homes and businesses were, on average, a thousand households away from a fiber node; now they are a hundred households away.

However, the imperial advance of this technology conceals a dark secret, which has led to a pervasive underestimation of the long term impact of photonics. Sixty percent of the fiber remains "dark" (unused for communications) and even the leading edge "lit" fiber is being used at less than one ten thousandths of its intrinsic capacity. This problem has prompted leaders in the industry, from Bill Gates and Andy Grove to Bob Metcalfe and Mitch Kapor, to underrate drastically the impact of fiber optics.

Restricting the speed and cost effectiveness of fiber has been an electronic bottleneck and a regulatory noose. In order for the signal to be amplified, regenerated, or switched, the light pulses had to be transformed into electronic pulses by optoelectronic converters. For all the talk of the speed of light, fiber optic systems therefore could pass bits no faster than the switching speed of transistors, which tops out at a cycle time of between 2.5 and 10 gigahertz. Telecom companies meanwhile could not deploy new low cost fiber products any faster than the switching speed of politicians and regulators, which tops out roughly at a cycle time of between 2.5 years and a rate of evolution measurable only by means of carbon 14.

Nonetheless, the intrinsic capacity of every fiber line is not 2.5 gigahertz. Nor is it even 25 gigahertz, which is roughly the capacity of all the frequencies commonly used in the air, from AM radio to Ka band satellite. The intrinsic capacity of every fiber thread, as thin as a human hair, is at the least one thousand times the capacity of what we call the "air". One thread could carry all the calls in America on the peak moment of Mother's day. One fiber thread could carry 25 times more bits than last year's average traffic load of all the world's communications networks put together: an estimated terabit (trillion bits) a second.

Over the last five years, technological breakthroughs and legislative loopholes have begun to open up this immense capacity to possible use. Following concepts pioneered and patented by David Payne at the University of Southhampton in England, a Bell Laboratories group led by Emmanuel Desurvire and Randy Giles developed a workable all optical device. They showed that a short stretch of fiber doped with erbium, a rare earth mineral, and excited by a cheap laser diode, can function as a powerful amplifier over fully 4.5 thousand gigahertz of the 25 thousand gigahertz span. Introduced by **Pirelli** (PIREY) of Italy and popularized by **Ciena** Corporation of Savage, Maryland and by **Lucent** (LU) and **Alcatel** (ALA), today such photonic amplifiers are a practical reality. Put in packages between two and three cubic inches in size, the erbium doped fiber amplifiers (EDFAs), fit anywhere in an optical network for enhancing signals without electronics.

This invention overcame the most fundamental disadvantage of optical networks compared to electronic networks. You can tap into an electronic network as often as desired without eroding the voltage signal. Although resistance and capacitance will leach away the current, there are no splitting losses in a voltage divider. Photonic signals, by contrast, suffer splitting losses every time they are tapped; they lose photons until eventually there are none left. The cheap and compact all optical amplifier solves this problem. It is an invention comparable in importance to the integrated circuit (IC).

Just as the IC made it possible to put an entire computer system on a single sliver of silicon, the all optical amplifier makes it possible to put an entire communications system on a seamless seine of silica–glass. Unleashing the law of the telecosm, it makes possible a new global economy of bandwidth abundance.

Five years ago when I first celebrated the radical implications of erbium doped amplifiers, skepticism reigned. I was summoned to Bellcore, where the first optical networks had been built and then abandoned, to learn the acute limits of the technology from Charles Brackett and his team. I had offered the vision of a broadband fibersphere–a worldwide web of glass and light–where computer users could tune into favored frequencies as readily as radios tune into frequencies in the atmosphere today. But Brackett and the other Bellcore experts told me that my basic assumption was false. It was no simpler, they explained, to tune into one of scores of frequencies on a fiber than to select time slots in a time division multiplexed (TDM) bitstream.

Indeed, electronic switching technology was moving faster than optical technology. In the face of the momentum and installed base of electronic switching and multiplexing, the fibersphere with hundreds of tunable frequencies would remain a fantasy, like Ted Nelson's Xanadu.

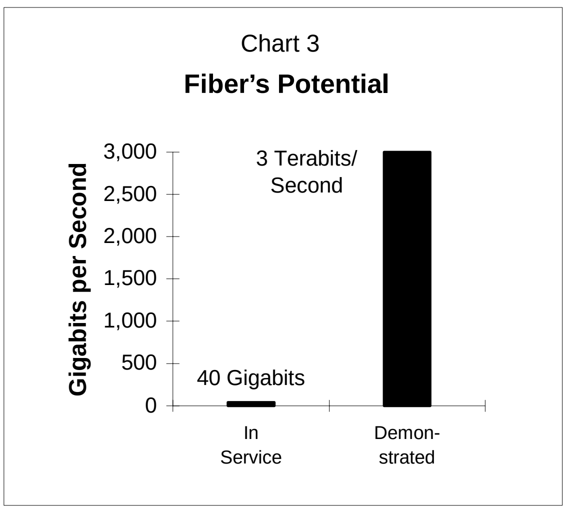

In 1997, the fantasy is coming true around the world. Xanadu has become the World Wide Web. The erbium doped fiber amplifier is an explosively growing $250 million business. Electronic TDM seems to have topped out at 2.5 gigabits a second, with OC 48. At OC 192 (10 Gb/s), TDM gear has suffered a series of delays and nagging defects and so far has failed in the market.

OC-192 TDM failed not only because it pushed the envelope of electronics but also because it violated the new paradigm. In single mode fiber, the two key impediments are nonlinearities in the glass and chromatic dispersion (the blurring of bit pulses because even in a single band different frequencies move at different speeds). Chromatic dispersion increases by the square of the bitrate, and the impact of nonlinearities rises with the power of the signal. High powered, high bitrate TDM flunked both telecosm tests.

By contrast, wavelength division multiplexing (WDM) follows the laws of the telecosm; it succeeds by wasting bandwidth and stinting on power. WDM takes some 33 percent more bandwidth per bit than TDM, but it reduces power to combat nonlinearity and divides the bitstream into multiple frequencies in order to combat dispersion. Thus it can extend the distance or increase capacity by a factor of four or more today and can lay the foundations for the fibersphere tomorrow.

In 1996, the new fiber paradigm emerged in full force. Parallel communications in all optical networks, long depicted as a broadband pipe dream, crushed all competitors and became the dominant source of new bandwidth in the world telecom network. The year began with a trifold explosion at the Fiber Optic Conference in San Jose when three companies–Lucent, **NTT** (NTT), and **Fujitsu** (FJTSY)–all announced terabit per second WDM transmissions down a single fiber. **Sprint Corporation** (FON) confirmed the significance of the laboratory breakthroughs by announcing deployment of Ciena's MultiWave 1600 WDM system, so called because it can increase the capacity of a single fiber thread by 1600 percent. It passes 16 data streams of 2.4 gigabits per second down a single fiber thread for 120 kilometers (compared with 35 kilometers and a single datastream for TDM).

The revolution continues in 1997. At the beginning of January, **NEC** (NIPNY) declared that by increasing the number of bits per hertz from one to three it had raised the laboratory WDM record to three terabits per second. During 1996, **MCI** (MCIC) had increased the speed of its Internet backbone by a factor of 25, from 45 megabits a second to 1.2 gigabits. On January 6, Fred Briggs, chief engineering officer at MCI, announced that his company is in the process of installing new WDM equipment from **Hitachi** (HIT) and Pirelli that increases the speed of its phone network backbone to 40 gigabits per second. Accelerating MCI's previous plans by some two years, the new system will use a more limited form of wavelength division multiplexing to put four 10 gigabit information streams on a single fiber thread.

The first deployment will use existing facilities on a 275 mile route between Chicago and St. Louis, but the technology will be extended to the entire network. This move will consummate a nearly thousandfold upgrade of the MCI backbone, from 45 megabits per second to 40 gigabits within some 36 months. Ciena meanwhile has announced technology that allows transmission of 100 gigabits per second.

About to go public in the most important IPO since Netscape, Ciena is the industry leader in open standard WDM gear. During the first six months

The intrinsic capacity of every fiber thread, as thin as a human hair, is at the least one thousand times the capacity of what we call "air".

US CHIPS MOVE FURTHER AHEAD

Chart 4

Semiconductor Market

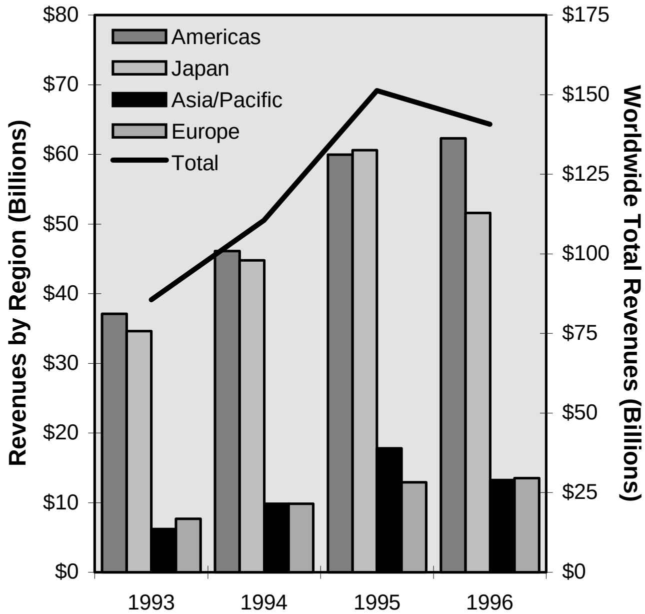

Source: Dataquest

Chart 5

Memory v Non-memory ICs

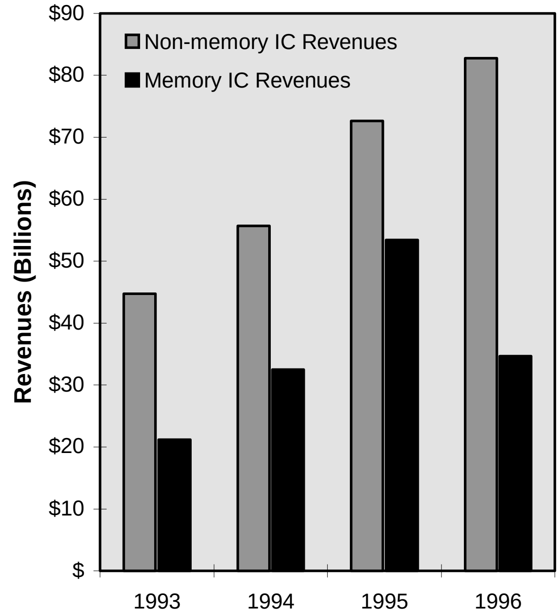

The seven percent 1995 to 1996 decline in the semiconductor market reported by Dataquest reflects the normalization of DRAM prices at the beginning of the year and the transition to higher density DRAM (see, GTR, October 1996). Consequently Japan and Asia with their relative dependence on DRAM suffered respective declines of 14.8% and 25.6%, while the Americas and Europe enjoyed growth of 3.9% and 4.9% respectively (Chart 4).

Chart 5 gives a preliminary breakdown separating memory integrated circuit (IC) revenues from non-memory ICs and illustrates the point—punctuated by early reports of record 4th quarter PC shipments—that falling DRAM revenues, rather than portending an industry-wide slowdown, can spur computer sales due to better price-performance (see, GTR, July 1996).

The geographic division of worldwide semiconductor and PC sales is shown in Chart 6 and Chart 7.

The worldwide distribution of computers in use—as distinct from sales—is presented in Chart 8. The number of computers in use for each region or group of countries is derived from Computer Industry Almanac statistics for individual countries for the end of 1995.

—KE

Chart 6

Worldwide Semiconductor Sales

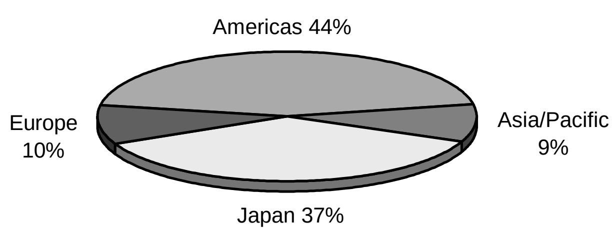

Chart 7

Worldwide PC Sales

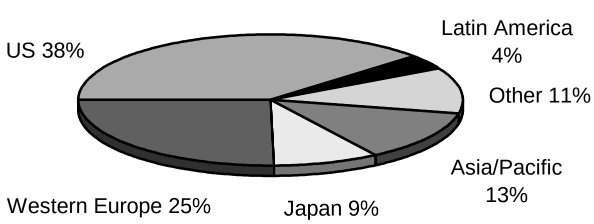

Chart 8

Worldwide Computer Penetration

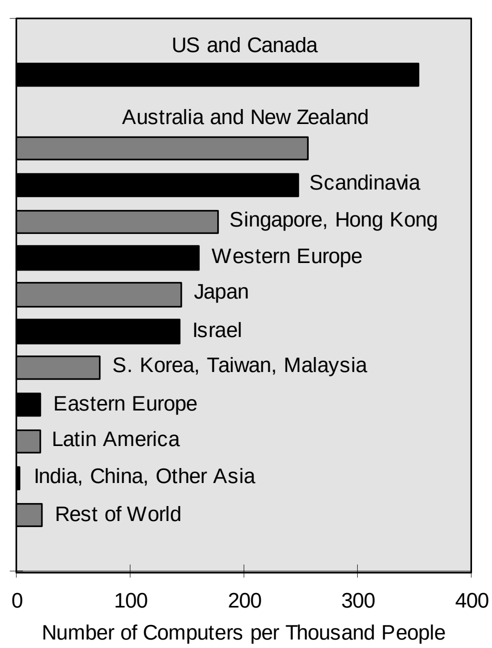

Chart 9

Internet Domains

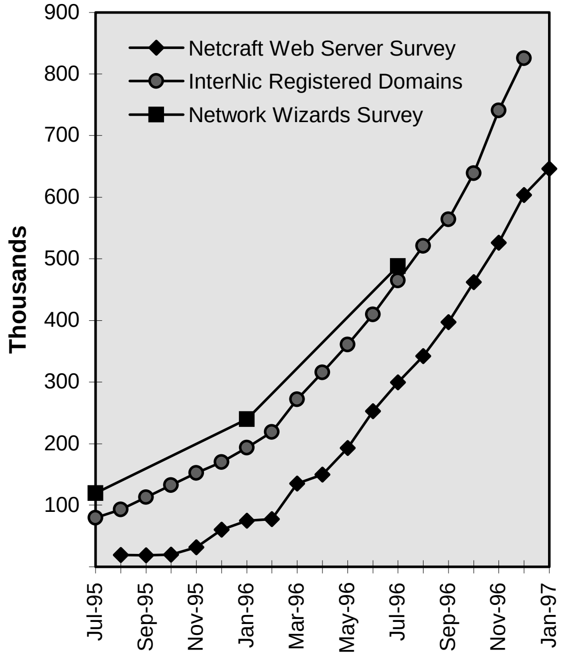

Chart 10

Internet Traffic

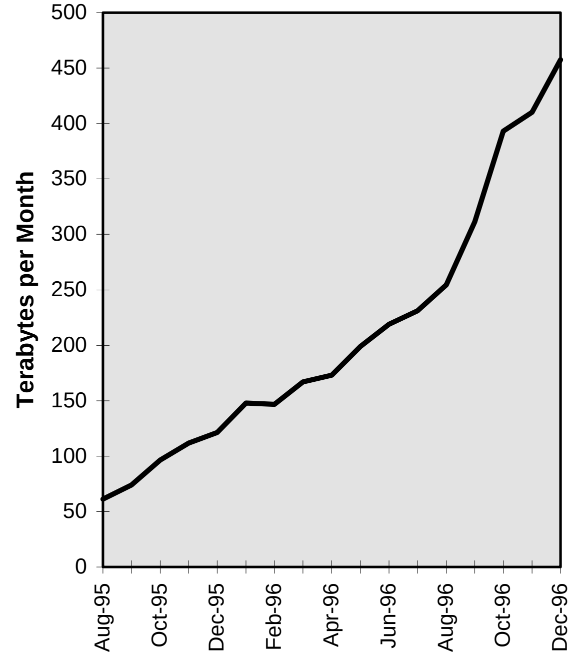

Despite the phenomenal growth of the Internet–385% increase in the number of registered domains from December 1, 1995 to 1996 (Chart 9) and 277% increase in traffic through the Network Access Points (NAPs) and Metropolitan Area Exchanges (MAEs) in 1996 (Chart 10)–the Internet remains healthy and unlikely to buckle under its own weight. Bob Metcalfe's predictions of certain collapse, which echoed throughout the popular media, were based in part on data collected by Merit Network, Inc. and its routing arbiter project. One year ago, Merit data with the inherently negative sounding names of delay, loss and instability were all climbing almost as fast as the Internet. Route servers at each of the five major exchange points recorded two sets of data for the routing arbiter project. Every 15 minutes the servers pinged each other across the Internet backbone and pinged other servers on networks connected through the NAPs. The time it took for the ping packets to go out and return was recorded as "delay" and the number of ping packets not returning determined "loss." (In this case, a lost packet is not analogous to lost email for the simple reason that ping packets have the lowest priority in Internet Protocol and are not retransmitted). The route servers also recorded data on the number of new routes announced and withdrawn, the combined total was labeled instability. While delay and lost data were presented as proof that the Internet was degrading in quality, instability was labeled as the big threat. With each route announced or withdrawn route tables across the Internet have to be updated. Excessive updates can–in the worst case scenario–lead to a route flap in which routers crash due to an endless loop of updates interrupting their primary job of routing. Unfortunately, the "bad" statistics have been repeatedly cited as evidence that the Internet is threatened. As recently as January 6, 1997, The EETimes led with the headline "Study: router flaws jamming Net traffic." The study cited was originally presented eight months earlier. Fortunately, the Merit data since May, 1996, (Charts 11 & 12) shows that delay, loss and instability all peaked last summer and have been stable and improving ever since in the face of large increases of net traffic.

–KE

Chart 11

Ping Packet Delay and Loss

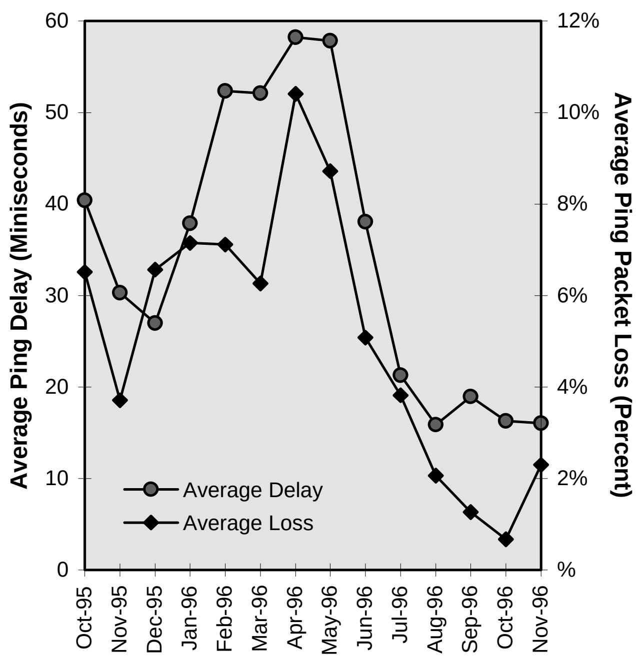

Chart 12

NAP/MAE Instability

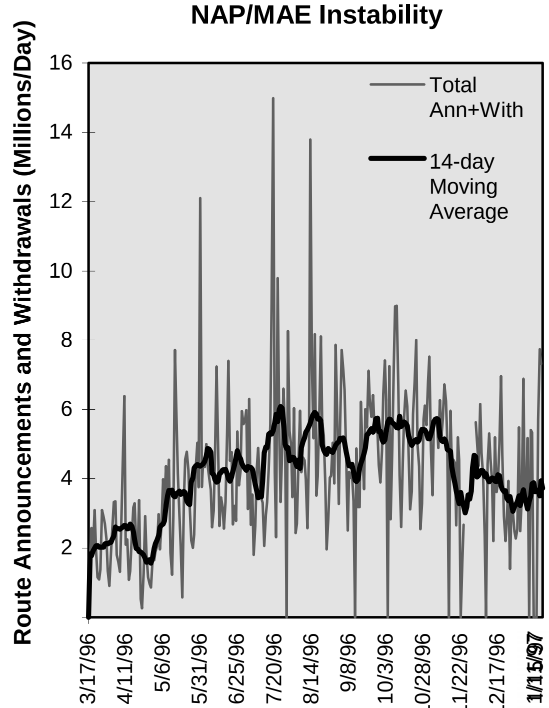

Chart 13

Companies on the Web

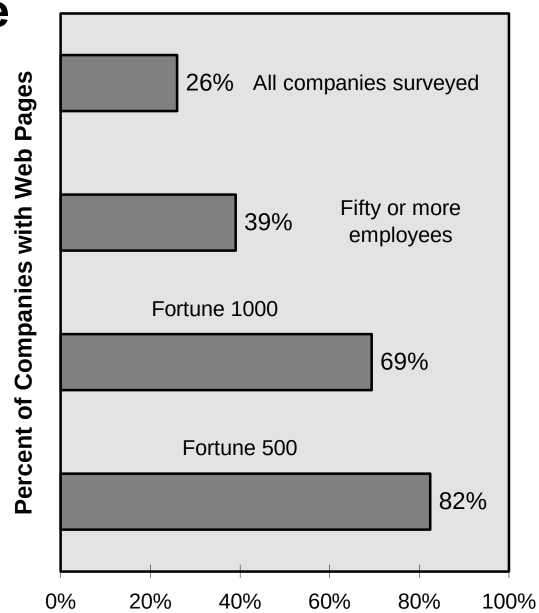

Chart 14

F1000 Web Server Software

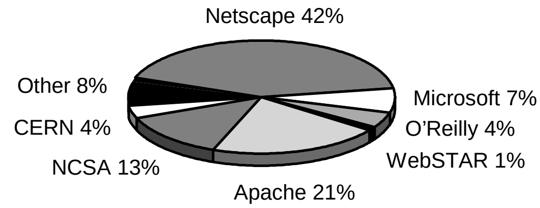

Twenty six percent of all companies have a web page according to a survey conducted by Odyssey for AT&T in the Fall of 1996. Among companies with 50 or more employees the percentage climbs to 39%. As we go to press, 69% of the top 1,000 companies and 82% of the top 500 have web sites (Chart 13).

Netscape is the web server software being used for 42% of the Fortune 1000 sites, while 38% have chosen public domain software from Apache, NCSA or CERN (Chart 14).

–KE

The fibersphere is a new paradigm and it will transform the architectures of telecom as deeply as the microprocessor paradigm revolutionized the architectures of computing.

the MultiWave 1600 was available, through October 1996, the firm achieved $54.8 million in sales and $20.8 million in net income. (Lucent is believed to be the overall leader with over $100 million of mostly proprietary AT&T systems). At the same time, the trans-Pacific consortium announced that it would deploy 100 gigabit per second WDM in its new link between the US and Asia.

The advance of the fibersphere is not restricted to long distance. Spearheaded by the Pentagon and seven private sector companies and led by Bellcore's Joseph Berthold, the $100 million "Monet" project seeks to move the fibersphere into the local exchange with an array of mostly passive optical multiplexing, splitting, and combining gear. Rod Anferness of Lucent, which is a Monet leader, reveals that optical add-drop multiplexers with fixed frequencies are selling in volume now. He estimates that tunable add-drop multiplexers, optical cross connect switches, wavelength converters, and other fibersphere components will come on the market before the turn to the century.

A powerful new player in these markets will be **Tellabs** (TLAB), currently the fastest growing supplier of electronic digital cross-connect switches and other TDM gear. In a further coup, following its purchase of broadband digital radio pioneer Steinbrecher, Tellabs has signed up all twelve principles in **IBM**'s all optical team. Headed by Paul Green, recent chairman of the IEEE Communications Society and author of the leading text on fiber networks and by Rajiv Ramaswami, co-author of a new 1997 text on the subject, the IBM group built the world's first fully functioning all optical networks (AONs), the Rainbow series. Tellabs now owns the 11 AON patents and 100 listed technology disclosures of the group.

More than mere writers, Green is the inventor of the rake receiver now widely used in CDMA telephony and Ramaswami helped him launch the "Muxmaster" (a disputed trademark, call it "Muxmaxer"), a 20 wavelength WDM product (IBM 9729). It links the trading floors of Morgan Stanley and other Wall Street firms to back offices in Brooklyn, Queens, and New Jersey over 20 one gigabit per second channels. Named a product of the year by Data Communications, the Muxmaxer brought IBM $30 million in revenues in 1996.

The implications of the WDM paradigm go beyond simple data pipes. The greatest impact of all optical technology will likely come in consumer markets. A portent is **Artel Video Systems** of Marlborough, MA, which recently introduced a fiber based WDM system that can transmit 48 digital video channels, 288 CD-quality audio bitstreams, and 64 data channels on one fiber line. Aggregating contributions from a variety of content sources–each on different fiber wavelengths–and delivering them to consumers who tune into favored frequencies on conventional cable, the Artel system represents a key step into the fibersphere. It can be used for new services by either cable TV companies or telcos.

The deeper significance of the Artel product, however, is its use of bandwidth as a replacement for transistors and switches. The Artel system works on dark fiber without compression. The video uses 200 megabit per second bitstreams (compare MPEG2 at 4 to 6 Mb/s), that permit lossless transmissions, suitable for medical imaging, and obviate dedicated processing of compression codes at the two ends.

A move to massively parallel communications analogous to the move to parallel computers, all optical networks promise nearly boundless bandwidth in fiber. According to Ewart Lowe of **British Telecom (**BTY**)**, whose labs at Martlesham

Heath in Ipswitch have been a fount of all optical technology, the new paradigm will reduce the cost of transport by a factor of 10. For example, the optoelectronic amplifiers previously used in fiber networks entailed nine power hungry bipolar microchips for each wavelength, rather than a simple loop of doped silica that covers scores of wavelengths.

As these systems move down through the network hierarchy, the growth of network bandwidth and cost effectiveness will not only outpace Moore's Law, it will also excel the rise in bandwidth within computers–their internal "buses" connecting their microprocessors to memory and input-output.

While MCI and Sprint move to deploy technology that functions at 40 gigabits a second, current computers and workstations command buses that run at a rate of close to one gigabit a second. This change in the relationship between the bandwidth of networks and the bandwidth of computers will transform the architecture of information technology. As Robert Lucky of Bellcore puts it, "Many of us have been conditioned to think that transmission is inherently expensive; that we should use switching and processing wherever possible to minimize transmission." This is the law of the microcosm. But as Lucky speculates, "The limitless bandwidth of fiber optics changes these assumptions. Perhaps we should transmit signals thousands of miles to avoid even the simplest processing function."

Lucky implies that the law of the telecosm eclipses the law of the microcosm. Actually, the law of the microcosm makes distributed computers (smart terminals) more efficient regardless of the cost of linking them together. The law of the telecosm makes broadband networks more efficient regardless of how numerous and smart are the terminals. Working together, however, these two laws of wires and switches impel ever more widely distributed information systems, with processing and memory in the optimal locations.

For the telephone companies, the age of ever smarter terminals mandates the emergence of ever dumber networks. Telephone companies may complain of the large costs of the transformation of their system, but they command capital budgets as large as the total revenues of the cable industry. Telcos may recoil in horror at the idea of dark fiber, but they command webs of the stuff ten times larger than any other industry. Dumb and dark networks may not fit the phone company self-image or advertising posture. But they promise larger markets than the current phone company plan to choke off their future in the labyrinthine nets of an "intelligent switching fabric" always behind schedule and full of software bugs.

Telephone switches (now 80 percent software) are already too complex to keep pace with the efflorescence of the Internet. While computers become ever more lean and mean, turning to reduced instruction set processors and Java stations, networks need to adopt reduced instruction set architectures. The ultimate in dumb and dark is the fibersphere now incubating in their magnificent laboratories.

The entrepreneurial folk in the computer industry may view this wrenching phone company adjustment with some satisfaction. But computer firms must also adjust. Now addicted to the use of transistors to solve the problems of limited bandwidth, the computer industry must use transistors to exploit the opportunities of nearly unlimited bandwidth. When home-based machines are optimized for manipulating high resolution digital video at high speeds, they will necessarily command what are now called supercomputer powers. This will mean that the dominant computer technology will emerge first not in the office market but in the consumer market. The major challenge for the computer industry is to change its focus from a few hundred million offices already full of computer technology to a billion living rooms now nearly devoid of it.

Cable companies possess the advantage of already owning dumb networks based on the essentials of the all optical model of broadcast and select–of customers seeking wavelengths or frequencies rather than switching circuits. Cable companies already provide all the programs to all the terminals and allow them to tune in to the desired messages. But the cable industry cannot become a full service supplier of telecommunications unless the regulators give up their ridiculous two-wire dream in which everyone competes with cable and no one makes any money. Cash poor and bandwidth rich, cable companies need to collaborate with telcos–which are cash rich and bandwidth poor–in a joint effort to create broadband systems in their own regions.

The law of the telecosm merges with the law of the microcosm. Electronic transistors use electrons to control, amplify, or switch electrons. But photonics differ radically from electronics. Because moving photons do not affect one another on contact, they cannot readily be used to control, amplify, or switch each other. For computing, photons are far inferior to electrons. For the foreseeable future, computers will be made with electrons.

Chart 15
N. American Cabled Optical Fiber
(1995)

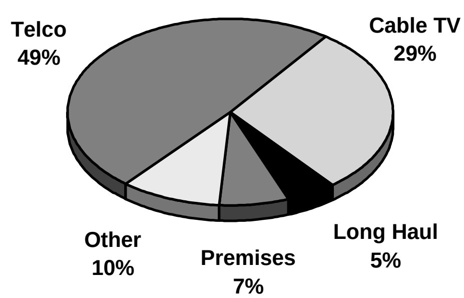

About to go public in the most important IPO since Netscape, Ciena is an industry leader.

CDMA CONTINUES GAINS

Chart 16

PCS Companies

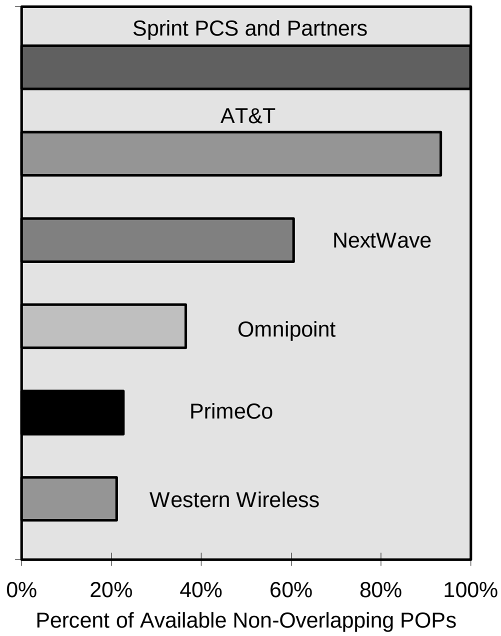

Chart 17

PCS Technology Commitments

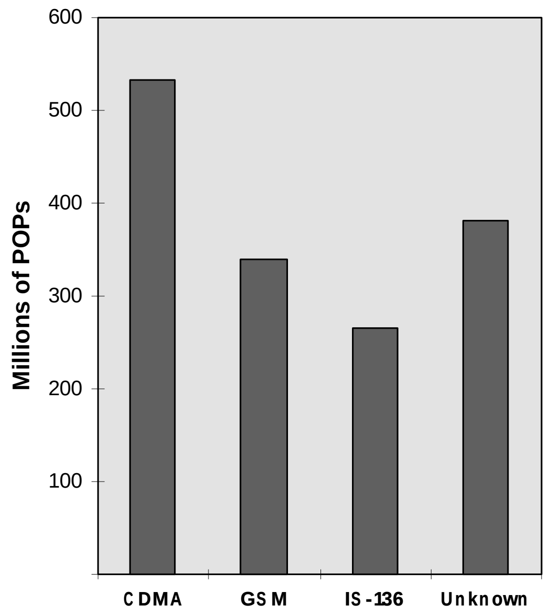

With the close of the D, E and F block auctions in January, 1997, the FCC has completed the auction of licenses to companies planning to offer Personal Communications Services (PCS). Each license grants a company the right to use one of six frequency blocks to offer services to the POPs (population) within a geographic region (Major or Basic Trading Area; MTA or BTA). As a result of the latest auctions, Sprint PCS (a consortium of Sprint, TCI, Cox and Comcast) will be able to offer its Sprint Spectrum CDMA based PCS service nationwide through its licenses, those of its partners Cox and Comcast, or American Personal Communications (49% owed by Sprint PCS). AT&T has expanded its licenses to cover 93% of all POPs, with agreements likely to extend its IS-136 (TDMA) based service to full coverage (if AT&T does not switch to CDMA). CDMA based NextWave and PrimeCo and GSM based Omnipoint and Western Wireless round out the list of the six companies with the broadest coverage (Chart 16).

With 200 odd companies purchasing licenses and countless alliances and roaming agreements still to be worked out, it is impossible to say what the final PCS landscape will look like, but a clear winner will be CDMA. Of the six potential service choices available to any subscriber (POP) within each of the BTAs or MTAs, two to three will be CDMA services (including Sprint PCS as one choice), one or two will be GSM services, and one will be IS-136 from AT&T. Many of the smaller license winners have either not yet chosen a technology or are negotiating with larger carriers wishing to expand service or fill in coverage gaps and their choice is dependent on future alliances. (Chart 17).

CDMA also advanced on the cellular front with Bell Atlantic NYNEX Mobile's January 21, 1997, announcement of CDMA digital cellular service in Washington/ Baltimore, Pittsburgh, and Charlotte, NC, with planned expansion to all its remaining markets by mid-year.

–KE

What are crippling flaws for photonic computing, however, are huge assets for communicating. Because moving photons do not collide with each other or respond to electronic charges, they are inherently a two way medium. They are immune to lightning strikes, electromagnetic pulses, or electrical power surges that destroy electronic equipment. Virtually noiseless and massless pulses of radiation, they move as fast and silently as light.

A natural division of labor cleaves photonics and electronics. Photonics will dominate communications and electronics will dominate computing. The two technologies do not compete; they are beautiful complements of each other.

In all eras, companies tend to prevail by maximizing the use of the cheapest resources. In the age of the fibersphere, they will use the huge intrinsic bandwidth of fiber, all 25 thousand gigahertz or more, to simplify everything else. This means replacing nearly all the hundreds of billions of dollars worth of switches, bridges, routers, converters, codecs, compressors, error correctors, and other devices, together with the trillions of lines of software code, that pervade the intelligent switching fabric of both telephone and computer networks.

The makers of all this equipment will resist mightily. But there is no chance that the old regime can prevail by fighting cheap and simple optics with costly and complex electronics and software.

The all optical network will triumph for the same reason that the integrated circuit triumphed: it is incomparably cheaper than the competition. Today, measured by the admittedly rough metric of MIPS per dollar, a personal computer is more than two thousand times more cost effective than a mainframe. Within 10 years, the all optical network will be thousands of times more cost effective than electronic networks. Just as the electron rules in computers, the photon will rule the waves of communication.

George Gilder, January 25, 1997

Gilder Technology Report is published by
Gilder Technology Group, Inc. and Forbes Inc.
Editor: George Gilder; Associate Editors: Charles Frank and David Minor;
Director of Research: Kenneth Ehrhart.

Monument Mills P.O. Box 660 Housatonic, MA 01236 USA
Tel: (413) 274-0211 Toll Free: (888) 484-2727 Fax: (413) 274-0213
Email: gtg@gilder.com