GILDER TECHNOLOGY REPORT

July 1996

Volume I Number 1

THE GILDER PARADIGM: “NOTHING BUT NET”

Amid the glut of bits and paper, *why do you need another newsletter*?

Amid the endless clink of lame coinages in the slot machines of the media, *why do you need another paradigm?*

You need it just to tame the glut and sort the onrush of new technologies, business concepts, investment hype, and PR hustles. Because I have been guided by a paradigm over the last several years, I have been able readily to dismiss many of the most touted new technologies, from **Zenith** HDTV and **Time-Warner** interactive television to **3DO** game machines and **Phillips** CD-I multimedia, from **McCaw** TDMA wireless and **Nextel** cellular compression to pervasive ATM (asynchronous transfer mode) networks. I could instantly embrace the developments that will be shaping the next decade, such as all optical networks, **Steinbrecher-Tellabs** smart radios, **Qualcomm** spread spectrum wireless, **Stratacom** frame relay, **MicroUnity** and **Chromatics** mediaprocessors, **Ipsilon** and **Netstar** switches, **Tandem** ServerNets, **Netscape** browsers, and **Sun**’s Java programming language. My paradigm did not trigger a torrent of quarters, or predict quarterly earnings, but it allowed me to find pivotal trends amid a welter of noise.

With the help of this newsletter, I hope you will be able to do it too.

I also need a letter. I need it to discipline my “visionary” journalism with specific data, predictions, and principles. I need constantly to test my paradigm against the facts, the numbers, and the news. I learned much of what I know about the current era by writing a series of reports on the microchip industry in the early 1980s as semiconductors editor for the [Ben] *Rosen Electronics*

Today, communications technologies are unleashing the Internet as the definitive force of a new industrial era, rendering the CPU peripheral, and the net central.

INTERNET CRUSHES SKEPTICS

Chart 1

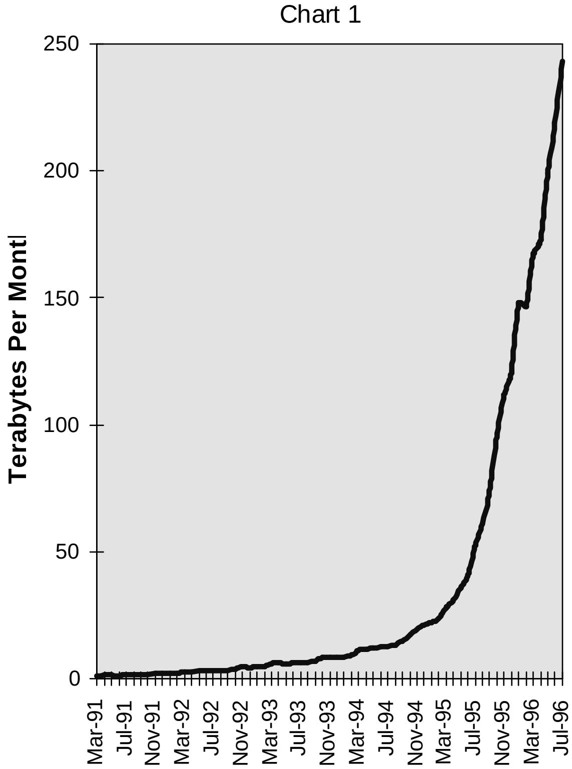

Estimates of the number of US Internet users range from 15 to 37 million. From one survey to the next the number may double or halve. InterNIC registered domain names have increased sevenfold from 50,000 in April of 1995 to over 360,000 in April 1996. But, these numbers include names being reserved for future use. Alta Vista and Lycos indexed web documents number from 22.7 to over 55 million respectively. What is the true story?

Chart 1 *shows the* ***growth of Internet use****, as measured in Terabytes of Internet traffic. Until now, the only statistics on traffic volume measured the Terabytes on the NSFnet backbone. In April 1995, that backbone and the corresponding data ceased to exist, following the transition to a commercial backbone interconnected at various Network Access Points (NAPs) and Metropolitan Area Exchanges (MAEs). We have calculated the traffic volume through these exchanges to bring the NSFnet statistics up to July 18,1996. This new data like the NSFnet data preceding it is incomplete. It does not include traffic transiting within a single network (Netcom, Uunet, etc.), between networks at private exchange points (MCI-Sprint, etc.)and at other exchanges which are either overseas, new or not reporting data (over 50). Furthermore bits/bytes transiting through two or more NAPs/MAEs on a single trip may be over represented here. But that said, this the most accurate data available, conclusively demonstrating the enormous growth of the Internet.*

Note: The blip in the Spring of 1996 reflects reporting problems at MAE East where traffic flow was underreported.

Technological paradigms are neither artificial nor arbitrary; they are the governing force in the practical life of human societies and economies.

Letter as it evolved into *Release 1.0* under Esther Dyson. While the microprocessor unleashed the personal computer as the definitive product of an industrial era, I developed my sense of the microcosm and its laws. Governed by the inexorable unfolding of [Intel's Gordon] Moore's Law, ordaining a doubling of transistors on a leading edge microchip every 18 months, the CPU (central processing unit) was indeed the central force in the world economy, rendering all other devices as peripheral in comparison.

Today, communications technologies are unleashing the Internet as the definitive force of a new industrial era, rendering the CPU peripheral, and the net central. This shift is even more fundamental than the rise of the PC and I want to cover it as fundamentally as I covered the emergence of the microcosm. This newsletter will drive me to do this–to pursue the new regime as persistently as I pursued the microcosm of the last 15 years, to uncover laws of the telecosm as compelling as the laws of the semiconductor era that I treated in my book *Microcosm: The Quantum Era in Economics and Technology*. The new model will unleash a tide of wealth creation that is already sure to dwarf the some $400 billion created by the personal computer.

I will write most of the reports myself. However, I will be assisted by several statisticians and researchers, led by Ken Ehrhart, and I will be seeking other writers to contribute to this exploration. Part of the adventure will come in using the technologies of the web to fathom its meaning and internal logic. Part of the adventure will come from the entrepreneurial challenge of creating a new company in the telecosm. Leading this effort will be Charles Frank and David Minor, who conceived Gilder Technology Group and launched it in Housatonic, Massachusetts. But the heart of the adventure is the discovery and development of the laws of the telecosm–the ruling dynamics of the new paradigm.

Thomas Kuhn, the father of paradigm theory, defined paradigms as cyclical structures of knowledge that predictably succumb to new structures. Seeing history as a kind of academic debate ultimately governed by scientific fashions, Kuhn denied the cumulative character of scientific truth. His was a classic error of the academy, failing to grasp the necessary link between science and engineering. The cumulative facts of science manifest themselves in technology, powerful and practical machines that cumulatively and empirically validate the new paradigms and generate new science in the process. They reflect not fashion or ephemera but the profound and permanent truths of the universe.

The laws of the telecosm originate in the microcosm. The law of the microcosm can be summed up: Take any number "n" transistors and put them on a single sliver of silicon and you will get "n" squared performance and value. This is essentially Mead's Law, discovered by Caltech's Carver Mead in the early 1960s. It is based on the power-delay product in semiconductor electronics–the relationship between the heat dissipation and the speed of transistors. Since heat can destroy a semiconductor circuit, these two characteristics are functionally related. Most analysts believed that as ever-smaller transistors, ever more sensitive to heat, were jammed together on single slivers of silicon, the power delay product would deteriorate. In this model, dense chips would sizzle like a frying pan, leading to circuit melt-downs.

Limiting the density of devices possible on a chip, this fear drove **IBM**, **Fujitsu**, **Hitachi**, **Cray**, and others to focus on increasing the speed of single switches, such as tunnel diodes, Josephson junctions, high electron mobility transistors, heterojunctions, gallium arsenide devices, and other exotic components in the mainframe paradigm. Exploring the physics of semiconductors at the quantum level, however, Mead showed that the power delay product improved exponentially as the size of the transistors diminished. As transistors were miniaturized and jammed more closely together, they ran faster, cooler, cheaper, better. The less the space the more the room.

Thus emerged the paradigm that governed the industry from 1969 until now. Mead's Law pointed the way to Moore's Law, which ordained that the number of transistors on a device could double every 18 months. This finding is still the foundation of the age of the personal computer. Combined with Mead's Law which shows that the smaller transistors will run exponentially better and cheaper, this finding produces crucial corollaries. The price per bit of computer memory will drop 68 percent per year. By extension from the physics of silicon to the related physics of metal oxides, the price per bit of hard drive storage will drop at a similar pace. As displays join the semiconductor learning curve, they too will drop in price at the same rate. Derivatively, the cost-effectiveness of computer hardware will double at least every 18 months.

These vectors of advance allowed me–and others–to predict with complete confidence that PCs would increasingly displace mainframes, supercomputers, minicomputers, timesharing terminals, television sets, telephones, and other centralized systems. Companies that tried to resist the tide, from IBM and **Digital** to Fujitsu and Hitachi to **Thinking Machines** and Cray Research, from Zenith to Phillips to **Sony**, would lose ground to companies that embraced the PC, from **Intel** and **Compaq** to **Adobe** and **Microsoft**, and to the companies that rode the explosive advance in the number and density of microchips, from **Applied Materials** to **KLA**.

Successful companies would follow the essential rule of thrift in the microcosmic era: waste transistors. If you tried to optimize every transistor on a million transistor chip, you would never get a product out the door. Bet on Moore's Law to bail out imperfect systems, such as Windows 1.0 or the x86 Intel microprocessor instruction set. Driven by this centrifugal paradigm, the share of total computer power commanded by mainframes and other centralized systems dropped from nearly 100 percent in 1977 to under one percent in 1987. By 1995, the mainframe was essentially a peripheral of the PC–a server supplying it with needed storage or transactional I-O (input-output).

The momentum of the microcosm continues to mount. Personal computers not only dwarfed the sales of all other kinds of computers; in 1996's first quarter, they outsold TVs in units by a whopping 25 percent–5.73 million PCs to 4.6 million TVs (I had prematurely declared this event for 1995). After a record billion dollar profit in the first quarter, Intel projects 60 million Pentium sales in 1996. In American local telephone companies for the first time databits outnumbered voice bits (and databits were rising 10 times faster). My prophesies of the death of television and telephony seemed less outlandish day by day.

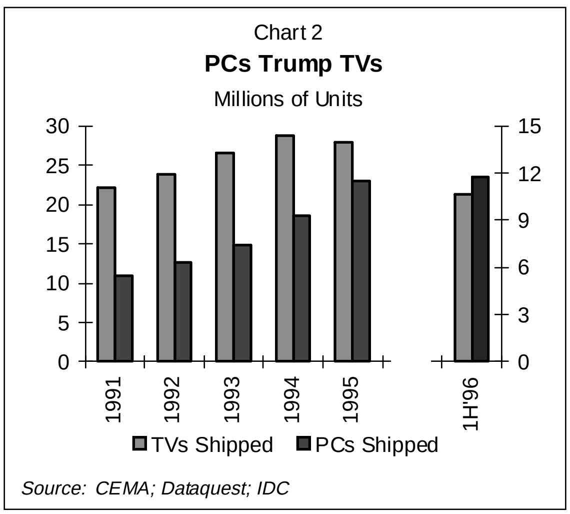

To grasp the microcosmic paradigm the key was to ask yourself a simple question: What would you do if transistors were free? The answer was: Put them everywhere.

In most ways, the power of the microcosmic paradigm to predict the future of microprocessors, computers, and their architectures has increased. If that was all there was to the story, this letter would not be necessary. But the PC paradigm is suffering problems. For example, software is widely seen to be in crisis. Ted Lewis, the seer-sage of *IEEE Computer* magazine, reports that while hardware in general advances at a pace of Moore's Law, around 48 percent per year (doubling every 18 months), microprocessors double their performance every 15 months (56 percent per year). But software functionality ekes up 4.5 percent per year, while sprawling imperiously across the exponentially growing space on hard drives.

Microsoft Office, for example, currently occupies some 50 megabytes; the Windows 95 operating system usurps 15 megabytes. Of course, these numbers probably mark a record low as a proportion of average disk space. The problem is that these programs are maladapted to the new machines. The programs depend on constant reads and writes to working memory. Yet for all the continuing surge of memory capacity, memory speeds are rising just 7 percent per year. Since the early 1980s, the ratio of CPU speed to memory speed has risen from essentially one-to-one to approach 20 to one, with CPU clock cycles at 3.5 nanoseconds and DRAM access times at some 60 nanoseconds. Operating software, independent of content, squanders the scarcest resource–the memory-processor link–at a time when this bus is also under increasing pressure from the network.

The chosen solution is to multiply memory and pile up caches of fast static RAM. But increasingly the bulk of content resides on the network. In the Wintel system, network access entails several further layers of operating system software and protocols, including unstable mazes of protocol stacks, comm ports and DLL files.

Microsoft is now trying to absorb all this complexity into a "thick client" operating system. Bill Gates still assumes that the old paradigm still prevails–that the desktop is the center and can adapt the network to its legacy systems. He sees Internet standards as a way for Microsoft to reach out and usurp the local area network in new Intranets based on the NT operating system. "The LAN is dead," declares Microsoft guru Nathan Myhrvold. Then the company is mincing out into a wider public arena in partnership with **NBC**, Hollywood, and Michael Kinsley–three more legacy systems lost in old paradigms–while Gates claims that the new paradigm of bandwidth abundance is "farther away than ever."

Lewis sums up this blend of Wintel film glitz, faster chips, interactive media, and code bloat as "Siliwood." With 80 percent of the desktop software market and an array of Internet ventures, Microsoft can afford its Siliwood games for some time. But the new paradigm will not roll over the LAN and the client server model without gaining the momentum to roll over the desktop as well. In the new paradigm, the installed base is your enemy. You cannot see the future through Windows.

Over the next two years, the Wintel model will scale its highest peaks of sales and earnings. Fifty to 75 percent lower prices for Pentiums, DRAMs, and portable displays will unleash a new boom before the transition. But the transition is inexorable. By early 1998, the new paradigm will blow away the current Wintel grip and open up a far

Chart 3

Pentium Price-Performance

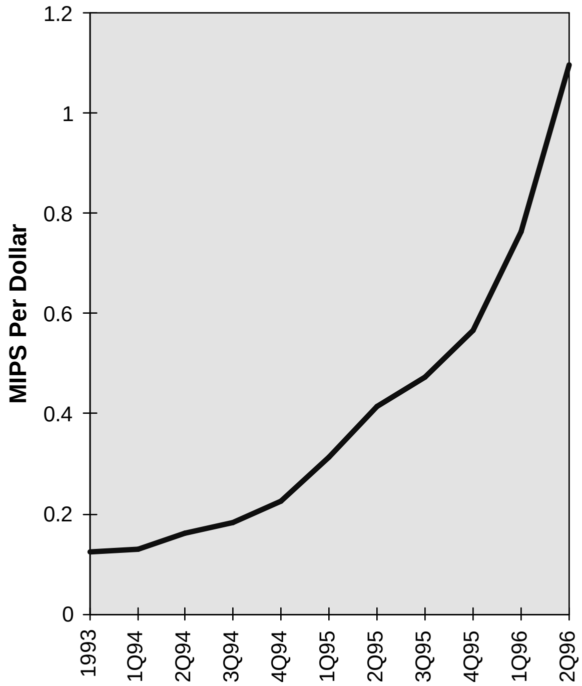

Chart 4

Bytes Per Dollar

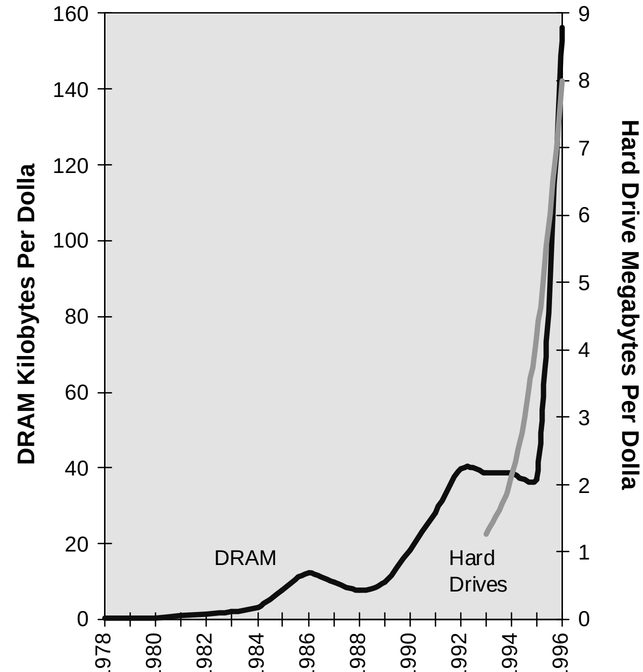

The practical importance of Moore's Law rests in the fact that a dramatic increase in the number of transistors on a chip leads to a dramatic drop in the cost of those transistors, resulting in a boost in price-performance. Intel has profited from these dynamics, using regular price cuts and performance improvements (Chart 3) to move the processor chip market from the 486 to the Pentium and to ward off clone makers. Hard drives have seen a similar boost in price-performance (Chart 4). But, the price-performance curve for DRAM was distorted in 1987-88 and again from 1992 to 1995 before the expected decline in bit prices resumed in 1996.

Meanwhile, software suites and multimedia content have been expanding to fill hard drives and use up memory, offsetting the cost benefits of hardware improvements. Chart 5 shows the dramatic increase in hard drive size and in the megabytes of DRAM produced (whether shipped in new PCs, used in upgrades or in peripherals) per PC shipped. Between 1991 and 1995, increasing use of DRAM while prices remained stagnant led to a near three-fold increase in the cost of DRAM produced per PC. The average size and cost of a hard drive also increased from 1993 to 1994 (Chart 6).

Fortunately, price-performance advances have resulted in a reversal of that trend. The four-fold increase in hard drive size from 1993 to 1996 has been essentially free because of the intervening drop in the cost of drives. And the 1996 improvement in DRAM prices has brought a great opportunity for PC buyers and producers. Instead of shipping an under-performing PC with 8 Megabytes of DRAM (at a 1995 DRAM cost of over $200) the manufacturer can now install 32 Megabytes of DRAM for under $200; or, a purchaser could be offered a middle-of-the-road 16 Megabyte machine at a savings of over $300!

*These dramatic DRAM, hard drive and processor component cost improvements, create the basis for strong computer sales. Factor in a four fold decline in flat panel display costs (**a 10" TFT display costing $1200 in 4th Q '94 cost $300 in 2nd Q '96**) which feed the thriving mobile computer market; strong international PC demand; along with untapped domestic opportunities for Network Computers and other Internet access devices and you have the basis for a computer boom.*

Chart 5

MegabytesPer PC

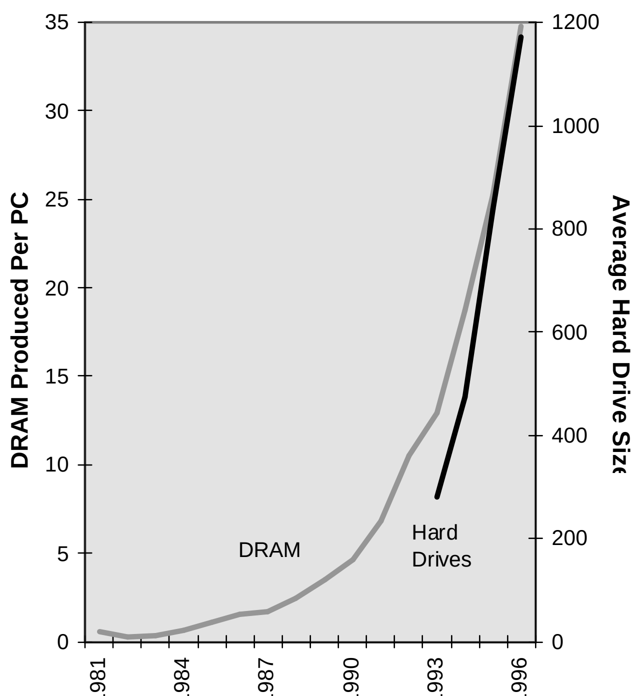

Chart 6

Storage Cost Per PC

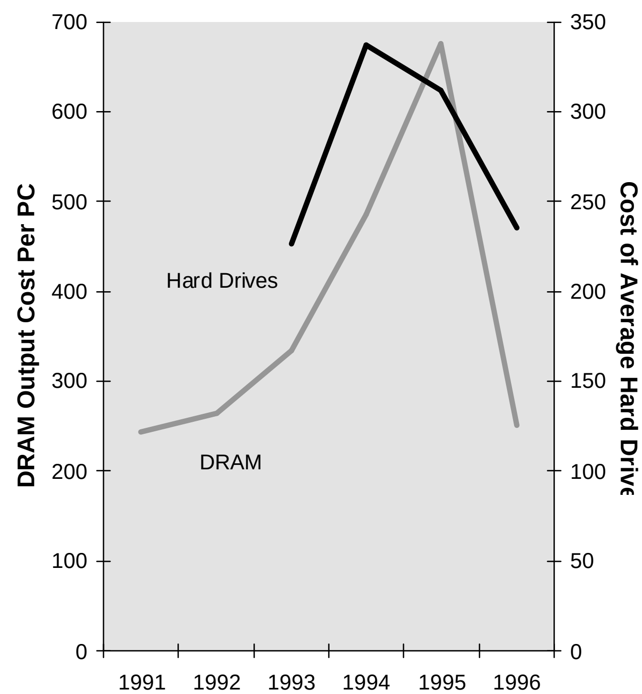

Chart 7

Web Pages With Java Applets

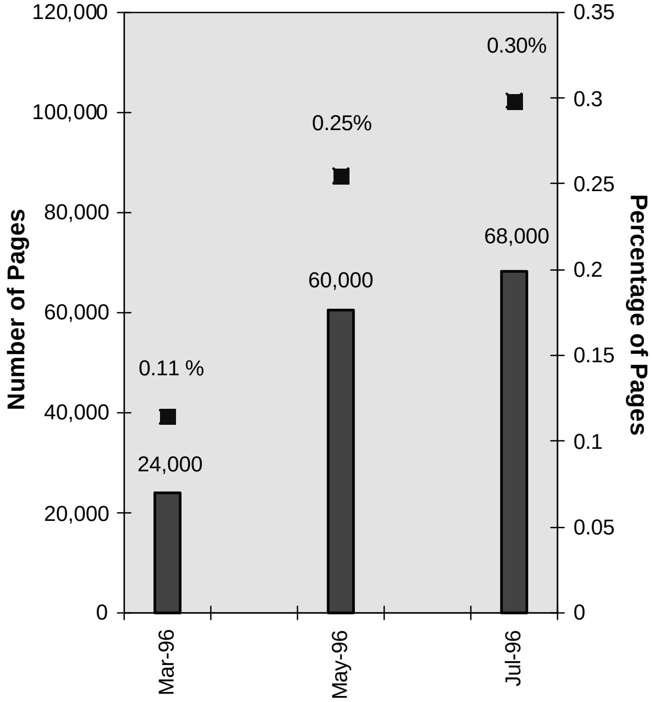

Java's use is increasing. *While over 68,000 Web documents with embedded Java Applets have been indexed by Alta Vista, progress is also being made toward full scale Java Applications, with over 200,000 lines of code. Lew Tucker, head of developer relations at JavaSoft, expects 200 to 300 full-size programs by the end of the year.*

Chart 8

Licensing of Java

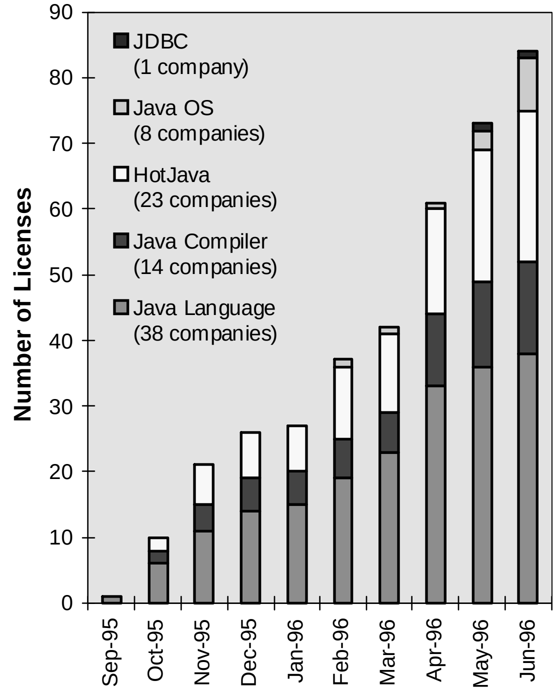

Business plans are increasingly included Java. *47 companies have licensed one or more Java component. 160 vendors displayed Java related wares at the JavaOne Developers Conference. Less than 50% of the companies JavaSoft is working with have made public statements, yet over 200 corporations have announced plans to build toward Java.*

100,000 Java Developers

Based on sales of "hard core" programming books, Mike Hendrickson of Addison-Wesley Publishing estimates the number of serious Java developers to be between 60,000 and 100,000. In comparison, hard core Windows developers (Visual Basic and Visual C++) number about 400,000, he says. 100,000 Java developers is also the "conservative" estimate of JavaSoft's Lew Tucker, based on the downloading of 90,000 Java Developers' Kits per month. These numbers will surely grow. In addition, to the over 200,000 people trying to teach themselves Java in 21 Days, there are at least 45 universities now teaching Java, according to Barbara Gordon, the Vice President of Academic and Research Computing at Sun Microsystems.

Chart 9

Java Book Sales

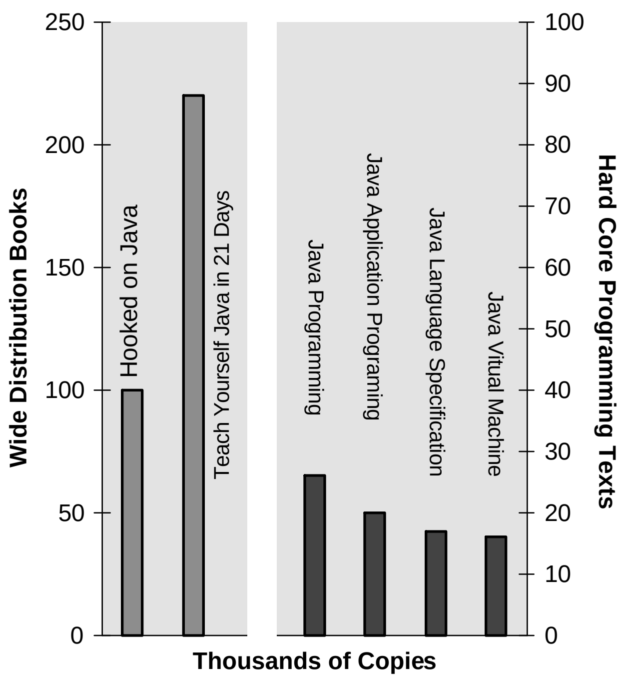

Chart 10

Number of Book Titles

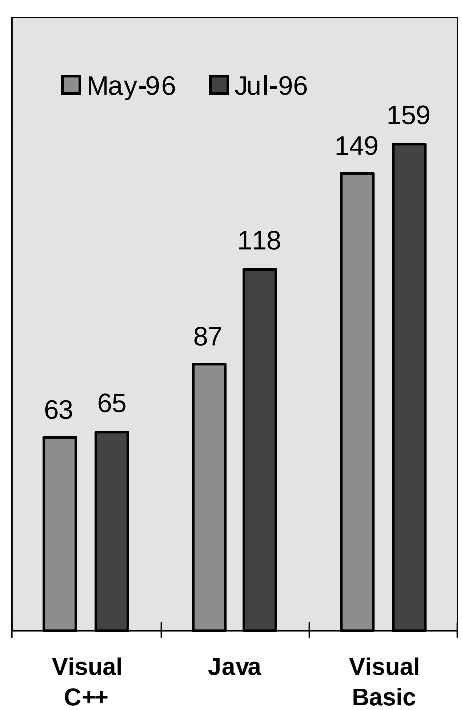

Companies Licensing Java

Netscape
Borland
Toshiba
Macromedia
Mitsubishi
Oracle
Metrowerks
Spyglass
Sun River
Adobe
SGI
IBM
AT&T
Symantec
Gupta
Microware
LSI Logic
Bandai
Hitachi
Microsoft
Novell
III Taiwan PC Consortium
Cybercash
NetManage
NEC
Pierian Springs
Justsystem Corporation
Dimension X
NCSA
SCO
Tandem
SunSoft
Apple
Parc-Place Digitalk
HP
Fujitsu
MagicWeb
Visigenic
Oracle
Wyse
Seiko
SME
Sybase
Toshiba
Acer
HDS
Crystal WebWorks
Corel
SMCC

A new paradigm is emerging that will transform the industry and the economy just as deeply as the Moore's Law paradigm did. It will eventually overthrow all the dominant technologies of the old paradigm.

more diverse and prosperous information economy. Shaped by the contrast between stilted DRAM prices and the still plummeting cost per bit of hard drive storage, the last five years have left us a lopsided PC (rich in drive, poor in DRAM) and a related skew in software (suites that sprawl across your disk and clog your working memory). By resuming the ultra-Moore cost-slope, falling DRAM prices will enable more balanced machines. The likely result is a sharp upside surprise in computer sales–and thus in semiconductors–through 1997.

With the prevailing computer architecture approaching a point of diminishing returns, an increasing share of these new PCs will be network machines, optimized for the coming invasion of Internet bandwidth, on land and sea and in the air. Huge operating systems and related applications, linked awkwardly to communications protocols, impose a rising burden on the memory-processor bus and fail to take advantage of the rise in network bandwidth. This problem offers a beckoning opportunity to launch new PC or teleputer architectures that are entirely optimized for the net.

Chart 11
DRAM Back on the Curve

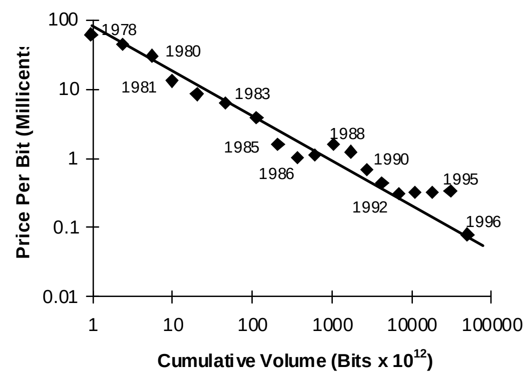

Source: ICE Status 96, Current Market Data

The microcosmic paradigm is giving way to the telecosmic paradigm–the law of the microcosm is giving up its supremacy to the law of the telecosm. On the critical datapaths of the new economy, peer networks of full duplex computers will roll over the stilted structures of client-server. Browser based systems with application specific devices adapted for Java and coupled to fast channels to the net will outperform the legacy systems from Microsoft and Intel. The datapaths to the future are on the net rather than the memory bus. This is not merely a trope of rhetoric, suggesting that networks are growing more important. Both Bill Gates and Andy Grove–the paragons of the old order–readily recognize the rising significance of networks. But both leaders challenge the speed and power of the change. Only partly in jest, Grove suggests a Grove's Law: Bandwidth grows 100 times more slowly than our power to use it. Gates is entranced by "middleband".

This letter takes a radically different position. A new paradigm is emerging that will transform the industry and the economy just as sharply as the Moore's Law paradigm did, when it wreaked a plunge of mainframe marketshare from 100 percent to below one percent between 1977 and 1987.

Since 1992, I have been developing this theory in *Forbes ASAP*, showing the journey toward the new order from many angles and several industries. But though the details are complex, the root point of view is startlingly simple. Every economic era is based on a key abundance and a key scarcity. In the pre-industrial era, horsepower was scarce and land relatively abundant. In the industrial era, horsepower (now known as watts) became abundant and land scarce, a vessel of resources. During the industrial revolution, "horsepower" plummeted to a cost of seven cents a kilowatt hour. Companies and industries gained market share by exploiting the kinetic energy released as this key factor of production hurtled over the new cliff of costs.

In the microcosmic era, transistors (ie mips and bits) joined with watts as abundant resources. To grasp the microcosmic paradigm, as Andrew Rapaport was first to observe, the key was to ask yourself a simple question: What would you do if transistors were free? The answer was: Put them everywhere. Squander redundant millions on every desktop. Bet against products based on expensive transistors, products that assume computing power will be rare and costly, software products that arduously economize on the use of memory, architectures that centralize intelligence and feed it to dumb client devices, whether phones, television sets, or desktop terminals.

As important to the paradigm shift as the abundance of mips and bits was the defining scarcity. What was scarce in the microcosm? One key scarcity was bandwidth–the communications power of wires and air. At favored frequencies, licenced users and regulations choked off the electromagnetic spectrum. At relevant prices, the public telephone network offered only four kilohertz copper wires. Satellites weighed more than a ton and offered a limited span of up to 225 megahertz of analog bandwidth feeding antennas a hundred feet in diameter. Thus, the modal economic activity of the old paradigm was to use transistors and watts as a replacement for wires and air. If transistors are free, you can use them to compensate for limited bandwidth by switching, routing, multiplexing, compressing, coding, buffering, and storing information.

Using watts as a replacement for bandwidth, you get radio and television stations sprawling over the air with high powered signals. You get cellular phone systems with one base station every thirty miles, with each radio tied to one exclusive span of spectrum. You get a world of computers and other appliances all constantly wired and plugged in.

Using switches as a substitute for bandwidth, you get the public switched telephone network, the 28.8 kilobit per second modem, the onset of narrowband ISDN, the persistence of analog cellular systems. You get half-rate vocoders for digital voice communications, compression technology as a Wall Street rage, CD-ROMs as a vessel for video inferior to NTSC TV. You get Andrew Grove keynoting the quadrennial International Telecommunications Union convention in Geneva and proudly showing off Proshare video teleconferencing at 15 frames per second. You get Microsoft selling you software suites full of programs you don't use but which use scores of megabytes of hard drive storage and 16 megabytes of DRAM, and giving away CD-ROM encyclopedias thin in data and full of dull animations. In a world of terabit per second transmissions in a single fiber thread, you get an entire worldwide communication net that carries a total of just one terabit per second. You get Bob Metcalfe and Howard Anderson predicting the collapse of the Internet because of a few terabytes per month of voice and video.

From telephone and computer networks to microprocessors and software, today's information economy is brimming with obsolescent architectures based on scarce bandwidth, free transistors, and free watts. The new paradigm will be based on the runaway expansion of bandwidth, outpacing Moore's Law by at least a factor of ten. But watts will become scarce, from lithium batteries feeding mobile computers, digital cellular phones, and undersea fiber amplifiers, and solar power fueling satellites. The bandwidth bottleneck will move from the network to the buses and I-O interfaces of the computer and no outpouring of watts will suffice to dissolve it.

With both wired and wireless bandwidth growing at least 10 times faster than computer power, the laws of the telecosm thus are eclipsing Moore's Law as the commanding force in the industry. To grasp the new era, you must imagine that bandwidth will be free and watts scarce. If the law of thrift in the old paradigm was *waste watts and transistors*, the law of thrift in the new paradigm will be *waste bandwidth and save watts.* In the new era, engineers will exploit the abundance of bandwidth and push the frontiers of low power technology to compensate for the limitations of existing computer and network architectures.

A precursor of the new computer architectures is Tandem's ServerNet system. Recently adopted by **NEC** and four other majors and being tested by IBM, it will be introduced by Compaq in the fall. Based on six-port router chips connecting "internal" computer functions across broadband network links rather than a shared bus, ServerNet brings the network into the computer and dissolves the distinction between I-O and interconnect. For many functions, such as video and transaction feeds to memory, ServerNet bypasses the CPU entirely to allow scalable broadband connections between memory, IO, and the internet. Extended to distributed computers, ServerNet becomes a multigigabit LAN technology suitable for direct access to large disk arrays and databases by hundreds of machines.

Other initiatives in the new paradigm for computer architecture are **MicroUnity**'s broadband mediaprocessors and **Chromatic**'s programmable broadband devices, which again offload bandwidth intense data from the CPU. Optimized for multigigabit links to the net and to I-O rather than for CPU instructions and processing, these programmable devices can handle such firehose gushes as video compression and decompression, 3 D texture reads and renders, echo cancellation, and microwave communications without the inflexibility of hardwired ASICs.

Nathan Myhrvold of Microsoft has observed that for high bandwidth communications an ATM link over fiber outperforms copper lines on a backplane bus. Extending that principle, the **Kendall Square Research** AllCache technology, being auctioned off on July 31 in Boston, offers a new unified global architecture, applying well known caching algorithms across a multiprocessor and onto a wide area network. Since 1988, Penn's gigabit guru David Farber has advocated a global computer based on similar shared non-uniform memory principles where the whole world is a cache. Farber's scheme dissolves the distinctions between I-O, network, and internal buses and "eliminates the field" of protocol design by conducting all communications through virtual shared memory accesses.

"Much to our surprise," Farber writes, "the software implications" for a megacomputer distributed across the globe "are essentially non-existent [since the system] is in all substantive respects identical to software...for a simple shared memory multiprocessor system." Virtual memory software similarly has to deal with drastically varying delay patterns from disk drives and RAM banks, which are effectively as far apart as Boston and Tokyo on a fiber web. All these approaches defer to the rising imperative of bandwidth as a substitute for processing, switching, and power.

If bandwidth is free, you get a completely different computer architecture and information economy. Transcending all previous

The new paradigm will be based on the runaway expansion of bandwidth… To grasp the new era, you must imagine that bandwidth will be free.

Chart 12

Fiber Deployed by Phone Companies

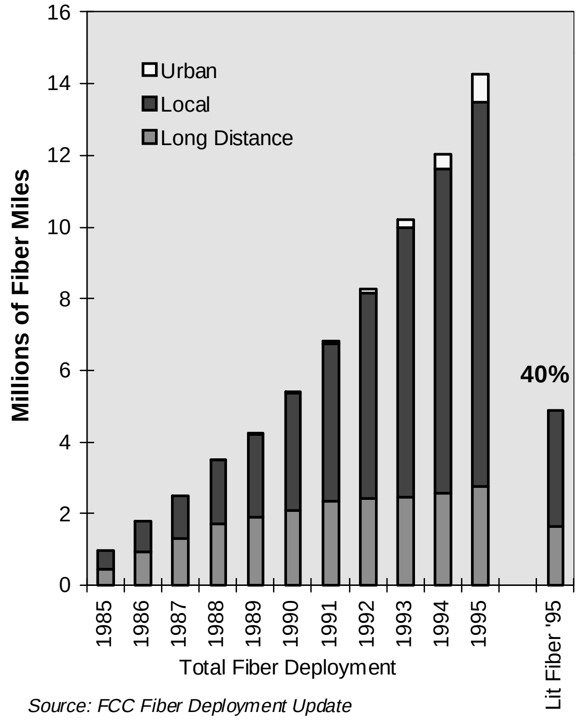

Source: FCC Fiber Deployment Update

Chart 13

Fiber Migration Toward the Home

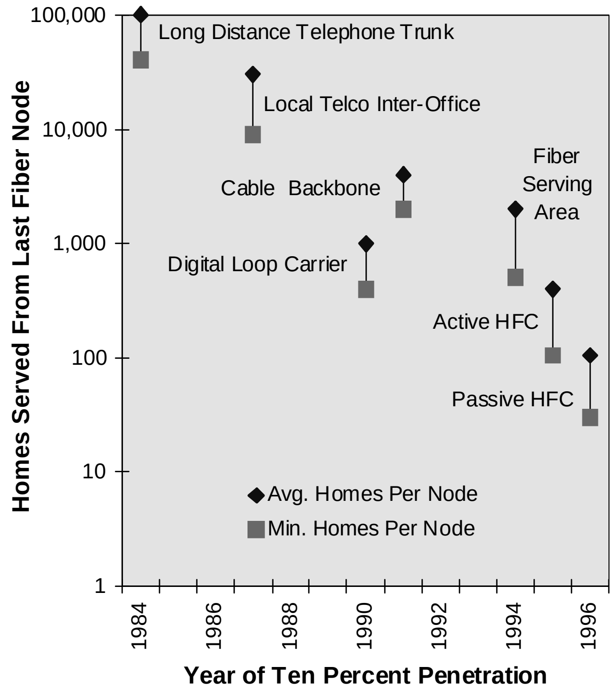

Source: D. Charlton, Corning

The transition from plain old telephone service (POTS) with its limited bandwidth capacity of 28.8 kilobits per second to the nearly unlimited bandwidth of fiber optics is a key to unleashing the power of the Telecosm. ***Chart 12*** *shows the deployment of fiber by phone companies over the past decade. The deployment by long distance carriers has slowed with their near total conversion to fiber. Deployment by local and urban carriers continues at a rapid pace as they bring fiber closer to the end user. Cable operators, meanwhile, have also been laying the fiber necessary to bring interactive services to their subscribers.*

Chart 13 *shows the penetration of fiber through the telecommunication infrastructure toward the home. Ten percent of the long distance telephone trunk had been converted to optical fiber by 1984. By 1987, fiber represented 10% of inter-office connections at the local phone company level. In 1990, Digital Loop Carriers converting analog signals coming over copper from homes to digital bits to be flashed through fiber back to phone company offices reached the 10% penetration level. In the cable industry, fiber was first used to connect the headend to remote satellite locations. It then penetrated across the cable backbone replacing existing coaxial trunk which had required numerous RF amplifiers along its path. By 1994 10% of homes were within Fiber Serving Areas in which an all-fiber trunk ends at a fiber node serving from 2,500 to 500 homes. Both phone and cable companies are now building Hybrid Fiber Coaxial networks (HFC) in which neighborhood fiber nodes connect with homes through a short coaxial line. Fiber to the Curb (FTTC) is a completely digital system in which fiber passes closer to homes, perhaps a street, block or apartment house (Fiber to the Basement) per node. The final leg might be coaxial, copper twisted pair or a completely wireless link. Fiber to the Home (FTTH) is exactly that ,the installation of a fiber connection on premise.*

Although projections suggest that FTTH will not reach 10% penetration until 2004, the dynamics are changing in favor of its earlier adoption. Problems remain in squeezing high bandwidth performance through thin copper wires and noisy coaxial links. The only "problem" with FTTH is a lingering belief that telephones should be powered by phone lines and the higher cost of laser and electronic components dedicated to each customer. Paul Shumate, Executive Director of Access Networks Technology Research at Bellcore, estimates that by moving the powering to customer premises with backup power supplied by batteries, the FTTH life cycle savings is about $200 per household. Furthermore, "maintenance and provisioning savings on metallic drops add another $200 savings over the 20-year lifespan," he noted. "The result is, there's really not much difference now and the numbers will only get better moving forward." The initial $150 to $400 premium paid for FTTH over the $1,000 to $1,100 cost of HFC or FTTC evaporates over time. With the increasing adoption of Fiber to the Desk (FTTD) within corporations and the start of full scale production of the electronic and optical components needed for FTTH we will see the arrival of the "free bandwidth" which will drive the economies of the Telecosm.

concepts of centralization and decentralization, one global machine distributes processing to the optimal point and access everywhere. Feeding on low power and high bandwidth, the most common computer of the new era will be a digital cellular phone with an IP address. **Geoworks**, **Nokia** and **Ericsson** are contriving some of its specs today. As mobile as your watch and as personal as your wallet, it will recognize speech, navigate streets, collect your mail and your paycheck, offer component software on demand, and serve up scores of thousands of choices for lifelong learning. It will link by infrared to a variety of displays, from tiny liquid crystal screens to giant micromirror projectors. It will connect through cellular broadband wireless systems to the Internet and the World Wide Web. The net will capture as much as one half of all the world's burgeoning commerce.

As early as 1949, Claude Shannon, the inventor of information theory, defined the crucial tradeoffs of a regime of bandwidth abundance. Bandwidth, he showed, can substitute both for switching and for power. The new paradigm requires that successful companies of the new era pursue this crucial trade off among the emerging technologies of sand and glass and air. *George Gilder – July 19,1996*

Gilder Technology Report *is published by* ***Gilder Technology Group, Inc.***
Monument Mills P.O. Box 660 Housatonic, MA 01236
Tel: (413) 274-0211 Toll Free: (888) GTG-2727 Fax: (413) 274-0213
Email: gtg@gilder.com