GILDER TECHNOLOGY REPORT

OCTOBER 1996 Volume I Number 3

CHECKMATE, JAVA

It was raining torrentially in Tokyo when I got my wakeup call. You may need one too.

It is 7:30 am and I am to speak within the hour at the cavernous new convention center near the new Rainbow suspension Bridge in the Airake section on the city's lavishly redeveloped waterfront. Looking out the window of the Nikkei Hotel at the gobs and slops of wind driven rain, I wonder if my speech will be canceled or delayed. As the car slogs through the flooded streets to the auditorium, I wonder whether anyone will show up. I shouldn't have worried. There would be an overflowing crowd of several thousand people at my speech and some 45 thousand would throng the center during the course of the morning.

Together with Java teleputers, Java promises to save a representative Forbes 500 company as much as $100 million a year.

I, however, am not the attraction. Nor is Hideo Nomo anywhere in view. These scores of thousands of Japanese are coming to learn a new computer language.

It is Java Day in Japan and no mere monsoon can thwart these crowds, representing all of Japan's leading and wannabe leading electronics firms, from learning this next new wave in computing. Java Day in Tokyo on July 10 ultimately attracted between 40,000 and 50,000 people to attend an array of formidable sessions on the intricacies of the language with Java's inventor James Gosling, key Java sponsor Eric Schmidt, and other **Sun** luminaries.

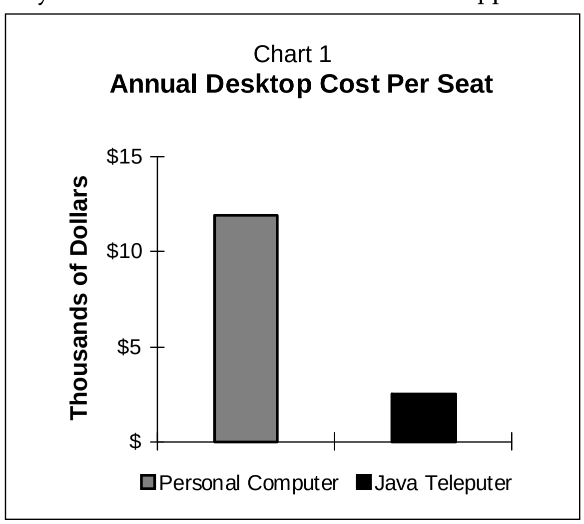

In Java's thrall, Japan is not unique. From San Jose where I visited the Embedded Hardware Processors' convention on September 17 to Berlin where I addressed the Open Software Foundation meeting on September 24, Java is everywhere ascendant. **Oracle** is creating an entire suite of industrial strength Java programs. At a conference in Monterey on October 3, I heard rumors that **Informix** and **Silicon Graphics** were leading some dozen other companies in a new Java Everywhere alliance. I learned that even in China there are some 20,000 programmers working on Java applications. Silicon Valley venture capitalists report that some 80 percent of the new business plans they see involve Java, and a consortium led by Kleiner-Perkins is launching a $100 million fund to support the technology.

Whether for network computing at the Webmaster Conference in San Francisco, Interactive CAD at Stanford's Center for Design Research, or high end computer architecture at MIT's Artificial Intelligence Lab, everywhere I go, whatever the question, the answer is Java. **EDS** executives enthusiastically confirm Sun's estimate that largely because of memory management and portability features, Java programmers excel C++ programmers by a factor of two or three in productivity.

Java is prevailing because in a multiplex world of proprietary systems optimized for various desktop processors, Java is shrewdly and resourcefully optimized for the net and for the network teleputers about to burst on the scene. Based on a generic Java interpreter, built into every new browser and operating system, Java is truly and

Within the next six to twelve months it will become obvious that the landscape of computing has been transformed—that the central standard in the industry is no longer Wintel but Java.

uniquely platform independent. With programmers writing for the generic machine rather than the proprietary OS, Java offers a compelling promise of write once, use anywhere software. It enables a system of component software rented and downloaded just-in-time from the net.

In a period when typical computer users spend more time accessing remote memories than local memories, it fits the new paradigm. It can save as much as $9,500 of the $11,900 that the **Gartner Group** estimates as the annual cost of an office PC. Together with Java teleputers, Java promises to save a representative Forbes 500 company, with 15 thousand available seats, as much as $100 million a year in hardware and software maintenance costs. Java systems render most PCs optimized for the desktop–however fast and fully featured–quite abruptly obsolete.

Java's new dominance will take most experts by surprise. Ted Lewis, guru in chief at IEEE Computer magazine, has been castigating Java for months as too slow and backward to be worth the considerable trouble of learning it. Bruce Eckel, author of books on both Java and C++, estimated in Web Techniques (Oct. 1996) that Java programs run between four and forty times as slowly as C++ programs of similar functionality. Sun responds with claims that this gap can be closed through the use of just-in-time compilers that adapt the code on the fly for particular hardware. But **Borland**, the concept's inventor, warned that Java would still run between five and ten times slower than C++.

Chart 2

Java Thin Client Savings

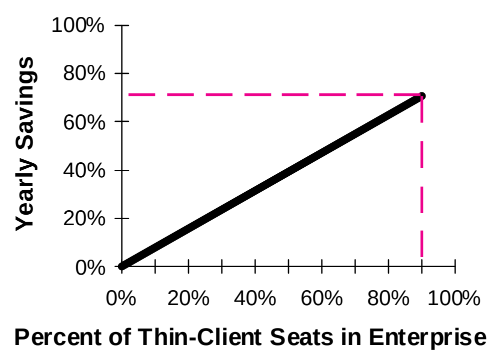

Meanwhile, **Microsoft** continues to gain share both in browsers and servers. By extending its lead in computer operating systems against **Apple**, **IBM**, and the various brands of Unix, Microsoft is triumphing even without winning any customers from **Netscape**. After all, Windows is prevailing in PCs and NT is advancing powerfully on high end corporate desktops and among ISPs, while IBM's OS2 still flounders and even **Hewlett Packard** is discarding Unix in favor of NT. While Sun will introduce a family of fast Java processors in early 1997, on the hardware side, quarter after quarter **Intel** roundly exceeds analysts expectations and technology projections.

Why, under these circumstances, should anyone want the interposition of Java's sticky interpretive middleware on top of their Microsoft OSs which run on 90 percent of the world's computers. Were there worms after all in James Gosling's Giant Peach? Many experts assured me that with some 400 thousand adepts and with 20 million desktops, Micro- soft's Visual Basic would dwarf Java in impact.

Then Microsoft mounted its relentless and resourceful campaign to absorb the Internet and Java in all its applications and operating systems while at the same time reserving precedence for its own object model, now termed Active-X, and its own browsers and servers. Long a Java skeptic ("webmasters would never learn it"), Stewart Alsop shook the industry with an astringent column in Fortune in early September declaring the likely decline and fall of Netscape, a key vessel of Java ubiquity. Microsoft even captured the most famous of the web computers, **WebTV**. In my travels, solemn experts told me that the entrepreneurial dreams of Mark Andreessen and Jim Clark had triumphed through hype and journalistic gullibility (no offense, George). Microsoft was already ascendant again and Netscape executives were discreetly selling shares at any legal opportunity.

Retreating to my hotel room TV in Berlin in late September, I heard Nathan Myhrvold of Microsoft tell Charlie Rose that "thin client" NCs–the crucial new vessels of Java in a network centric world–will never fly. With the ascendancy of graphics and video, people will need machines more powerful, not less. My imprudent prophecy at a speech to **Andersen Consulting** in June that Microsoft's market cap would never exceed its summer peak, based on a stock price of $112, seemed to rocket the Redmond giant to new highs. The shares stood at $133 when we went to press.

Oh, well, the higher they rise the harder they'll fall.

Don't look now, but Java has already won, and the Java computer is about to burst imperiously on the scene. Within the next six to twelve months it will become obvious that the landscape of computing has been transformed–that the central standard in the industry is no longer Wintel but Java. It will become clear that charging the blue flag flaunted by Marc Andreessen and Jim Clark, Microsoft overreacted to the Netscape challenge in browsers and servers and failed to meet the more profound crisis in computer architecture and software multiplexity.

When Windows 95 gave way to Windows NT as the favored next generation operating system for corporations, the game ended. Checkmate, Java. Sorry Bill. Subject to the same structural forces that afflicted the industry as a whole, Microsoft like IBM before it was incapable of maintaining a single standard. Ironically, NT, seemingly feared by Sun, assured Sun's ultimate ascendancy by creating a multiplex software environment, in need of Java's mediation, even in Microsoft itself.

By the end of the year, the Microsoft domain, like the world of Unix, would be riven by no fewer than five distinct operating systems–DOS, Windows 3.1, Windows 95, Windows NT, and a new Windows CE for consumer products. Collectively these operating systems comprise hundreds of thousands of applications, many orphaned by their creators, that could be ported only after recompilation and extended testing. From MVX to OS2, from AIX to AS400, IBM presented a comparable array of incompatible systems. For all the concern with open systems, Unix was similarly bursting out all over. The Internet was ineluctably the scene of rampant multiplexity, crying out in throes of fragmentation for the balms of Java. For all the focus of the industry on marketing trivia and market share shuffles, the tectonic plates of technological progress had shifted inexorably.

Wintel at 16

Just as teenagers rebel against Dad and Mom, leave home, and launch a new life, in the late 1990s, your computer is doing the same thing, rebelling against Pa Intel and Ma Microsoft. No hard feelings. The change in the industry has little or nothing to do with hostility to the 16 year old Wintel structure that still fuels industry growth. People still love their PCs. The change derives from a fundamental shift in computer architecture as the network becomes central and the CPU becomes peripheral. The result will be a new computer architecture,hardware and software,a Java-based network computer or teleputer that focuses not on displacing Dad and Mom but in functioning successfully in the world.

Chart 3

Processor-DRAM Perfomance Gap

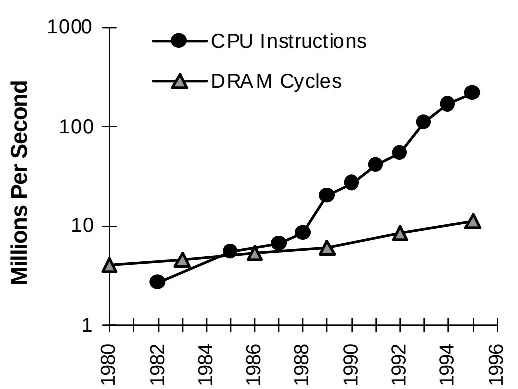

For thirty years, according to a calculation by Lewis in IEEE Computer magazine, hardware has been improving at a pace of 48 percent per year, while software advanced at a rate of 4.5 percent. This is no trivial difference; it's as if one leg of your teenager was growing ten times as fast as the other. Manifested in fatware or codebloat, the result, regardless of the heroics of Andy Grove, or Gerry Parker, the master of Intel's wafer fabs, is that software runs ever slower.

As Nicholas Negroponte of MIT puts it, Andy makes a faster chip and Bill takes all of it. Microsoft has not changed its ways. Its celebrated Explorer browser takes 76 percent more code than Netscape's 3.1, which has richer 3-D functionality and far greater portability. With comparable add-ons, Explorer is nearly twice as big. Under the old regime, codebloat and portability did not matter. But with the ascendancy of network computers and "Java thin clients" from Sun and others, portability and elegance will be decisive. With Java based executable content replacing many of the current "plug-ins," Netscape's leaner and more modular system will come into its own.

The crisis of computer architecture stems from six key forces. In each case, the Wintel structure is part of the problem rather than a part of the solution. It would fall of its own weight even if the Java computing model, based on thin client NCs, did not offer a powerful, multifaceted, and perfectly targeted solution.

1) Rising less than one tenth as fast as hardware speeds, software performance increasingly bogs down the machine. Yet Microsoft continues to launch ever larger suites that use ever more memory space and processor cycles.

2) Processor speed rises 60 percent per year while memory access improves just seven percent per year. This means, in effect, that not only the teenager's legs, but also his lungs and heart are growing at wildly different rates. The result is that processors, from **Digital** Alphas to Pentium Pros, spend between 75 and 90 percent of their time waiting on memory. The Wintel architecture offers no answers, no relief, and throngs the processor-memory bottleneck with ever growing traffic.

3) In the last four years, the percent of the delay budget devoted to interconnect latency in typical computer logic devices rose from 30 percent to 80 percent. The faster the switch, the longer it waits. Hurry up and wait is the Wintel mantra.

4) Over the last four years, the share of computer costs focused on local memory access through caches and other costly enhancements has risen at least 20 percent. But the Internet has reduced the percent of memory accesses focused on local memory from nearly 100 percent to under 50 percent.

5) Rather than the idiosyncratic array of Wintel instructions, what matters most in computing today is real time processing of digital signals, sounds and images, interfacing with networks, filtering, encrypting, decrypting data streams, and pattern matching. While Intel plays with putting MMX DSP code on the Pentium, DSPs in the guise of "mediaprocessors" will increasingly become the CPUs, integrated on memories, possibly DRAM.

6) Estimated by Gartner Group at $11,900, annual PC maintenance costs have risen to a level four times PC purchase prices. A key source of the problem is architectural multiplexity: the constant need to port and upgrade programs. Yet with a stream of new architectures, the Wintel establishment is steadily exacerbating the problem. A move to the Java paradigm could save literally hundreds of millions of dollars at Forbes 1000 companies. Checkmate, Java.

What could be wrong with the paradigm that has made Intel the world's largest chip firm?

DRAM AND CDMA ADVANCE

Chart 4
DRAM Chips and Bits

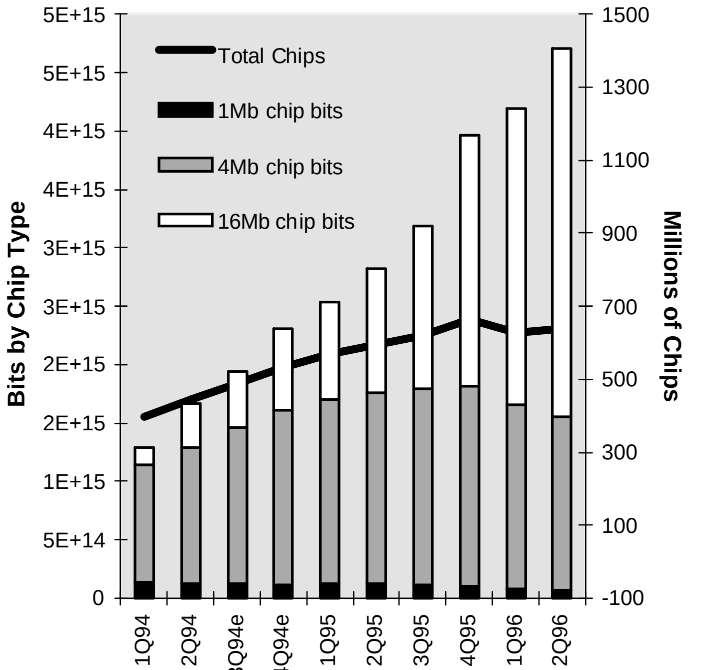

Chart 5
DRAM Fab Capacity

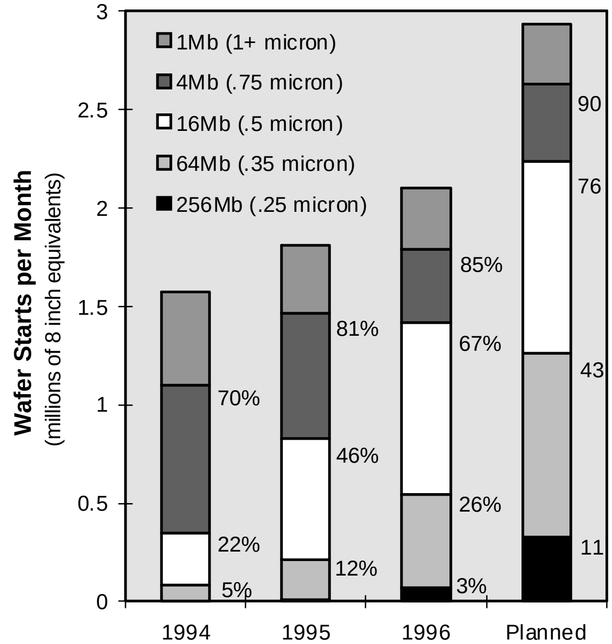

The decline in DRAM **chip** *sales from 4Q95 to 1Q96 and flat sales through 2Q96 (2Q96 up 7% from 2Q95) masks real advances in DRAM sales as measured in* **bits** *(2Q96 up 67% from 2Q95) (**Chart 4**). The defining fact during this period is that the previously dominant 4 Mbit chip has had its position usurped by the 4 times larger capacity 16 Mbit chip–fewer chips are required even while demand for bits continues to increase. This distinction must be understood to prevent false attempts to link flat DRAM* chip *sales to supposed stagnation in the PC or peripheral markets. Investment in fabrication capacity which uses ever smaller feature geometries allows both the production of denser smaller cheaper chips as well as higher capacity, next generation chips. In 1994, 70% of capacity was 0.75 micron or smaller, capable of producing 4 Mb chips, while only 22% was capable of producing 16 Mb chips. In 1996, nearly 70% of capacity is usable to produce 16 Mb chips or cheaper 4 Mb chips (**Chart 5**). These DRAM transitions and their market impact will continue to occur, with perhaps greater impact as the industry debates the leapfrog of 64 Mb to 256 Mb chips. The DRAM manufactures who are to succeed will continue capital spending to improve technology and lead the transitions or as Samsung is doing, will develop the technology to move beyond cookie-cutter DRAM to IRAM (Intelligent DRAM) incorporating processors onto DRAM and entire systems on single chips.*

*Even as Time Division Multiple Access (TDMA) based GSM digital wireless phone networks continue to gain subscribers worldwide (**Chart 6**), the Code Division Multiple Access (CDMA) standard continues gaining support in the US as the leading technology of choice among the winners of FCC licenses to offer advanced digital Personal Communications Services (PCS) (**Chart 7**). The FCC has so far auctioned 3 of 6 frequency blocks covering the US, which are composed of Major and Basic Trading Areas (MTAs and BTAs), each having associated with it a population number of potential subscribers (pops). The Sprint Telecommunications Venture (the alliance between Sprint and cable MSOs, including TCI, Cox, and Comcast) with 156 mil. pops, NextWave Communications (now in an alliance with MCI for the resale of service minutes) with 104 mil. pops and PCS PrimeCo (the alliance between RBOCs Bell Atlantic, NYNEX, US West and AirTouch) with 58 mil. pops, among others have all selected CDMA. Meanwhile, AT&T Wireless (107 mil. pops) has begun offering TDMA service based on the IS-136 standard.*

Chart 6
GSM Worldwide Growth

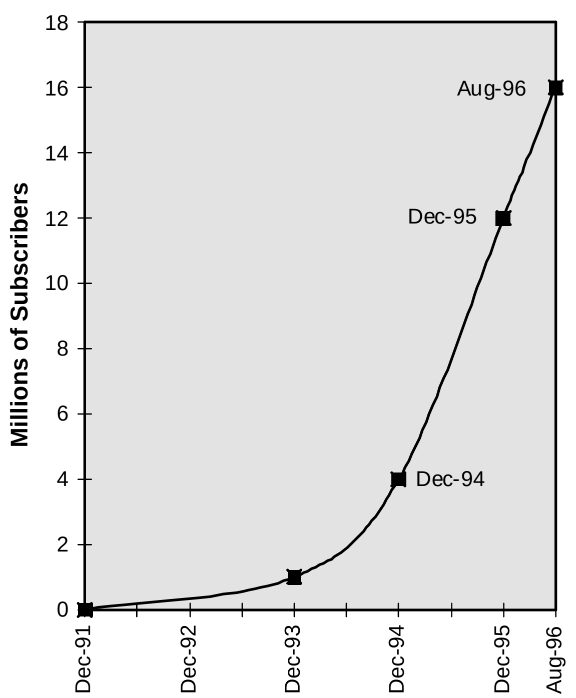

Chart 7
US PCS Technology Choices

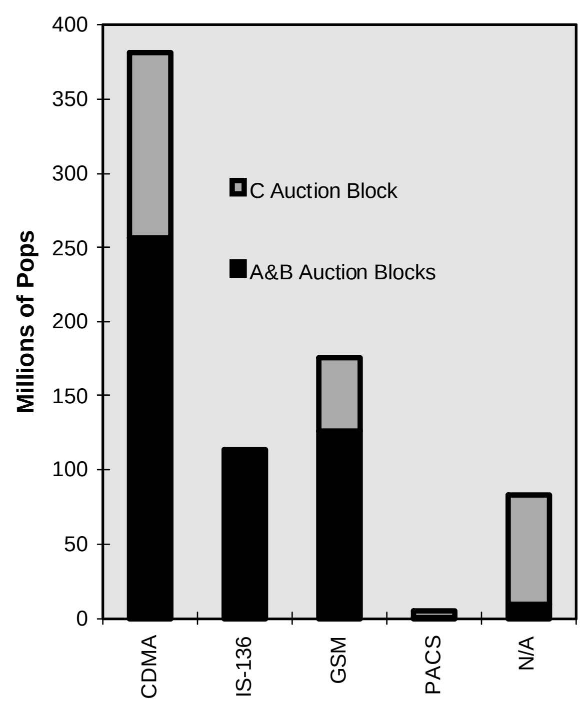

Chart 8

Web Pages With Java Applets

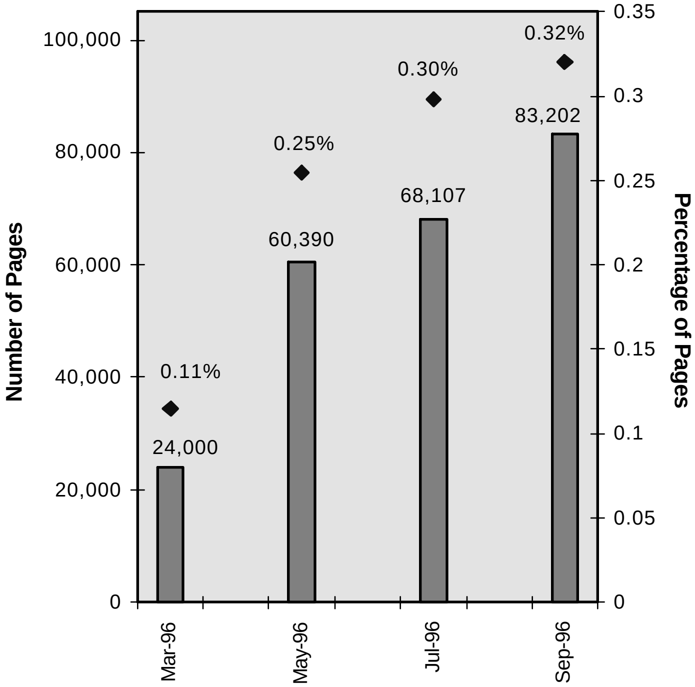

Chart 9

Java Book Sales

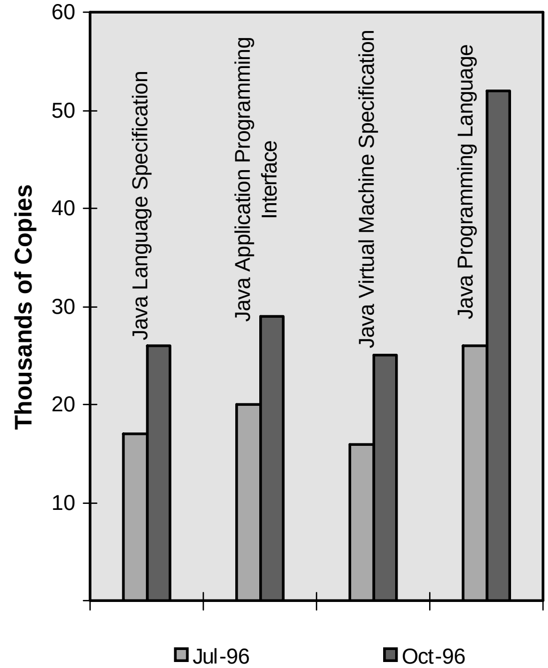

*The 247% increase in the number of web pages containing Java applets during the last six months (as indexed by Alta Vista) indicates the quick acceptance of Java by webmasters seeking to enliven web content (**Chart 8**). But, moving beyond simple applets, Sun (with Java Workshop written entirely in Java), Corel (with its office suite) and others are beginning to take Java to the next level of fully functional network ready applications. Wherever one looks Java is ascendant. Based on sales of "hard core" programming texts (**Chart 9**), Mike Hendrickson of Addison Wesley Longman publishing says the size of the market for Java books has clearly surpassed that of C++. In just three months, he has doubled his estimate of the number of "serious" Java programmers to as many as 200,000. But he stresses that his estimate may quickly rise to include all of the 400,000 or so "serious" Windows (currently Visual Basic and Visual C++) programmers. With Microsoft's embrace of Java within its operating systems and its release of Visual J++ for Java development, Java has the potential to quickly become the dominant language for all programming needs. Hendrickson points out that Java's features and the relative ease of Java programming are expanding the ranks of programmers. Paraphrasing, he suggests that whereas before only "A" or "B" students might have learned C++, even "C" students are learning Java. With over 57 colleges and universities now teaching Java, the potential is enormous.*

The size of the Internet, *whether measured by usage or infrastructure, continued its phenomenal growth in September. Internet traffic, flowing through the NAPs and MAEs, continuing the climb noted at the end of August, increased 22% from August to September (**Chart 10**). Last year's data, which showed an even larger 30% increase from September to October, and the early October data suggest the strong growth will continue this month. Meanwhile, the number of public web servers counted and polled by Netcraft (http://www.netcraft.com/survey) in their September 1st and October 1st surveys increased 16% **each** month, with Apache and Microsoft gaining the most marketshare (**Chart 11**).*

Chart 10

Internet Traffic

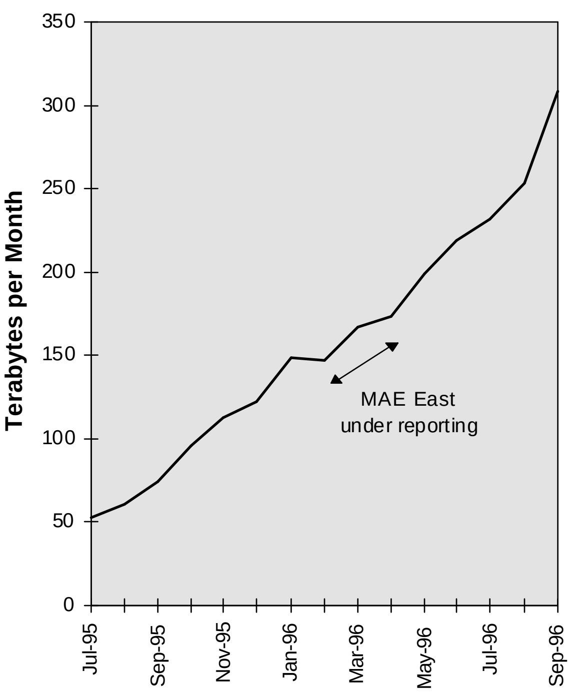

Chart 11

Internet Web Servers

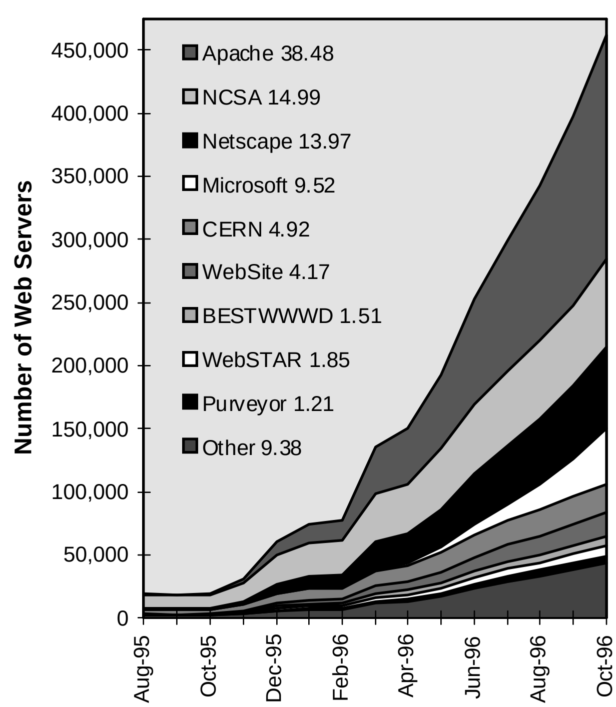

The fact is that Texas Instruments, the world DSP leader and the American DRAM leader, is today better oriented to the new era than Intel is.

Chart 12
Worldwide DSP Market

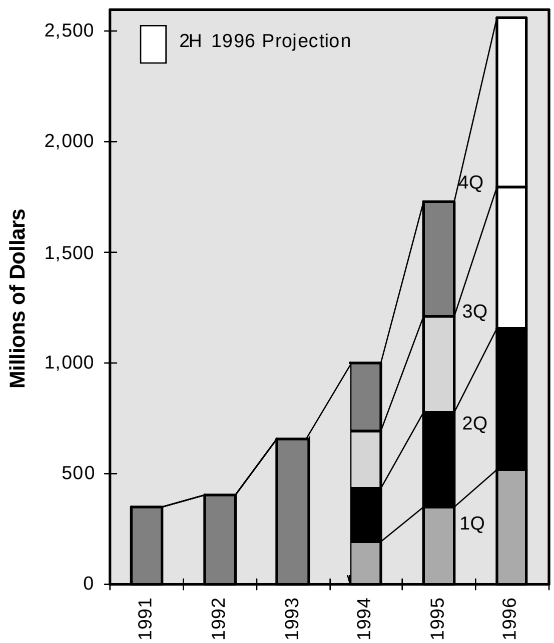

Chart 13
1995 DSP Market Share

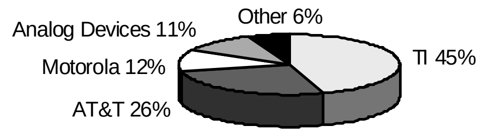

Source: ICE, "Status 96"

*Capable of real-time processing of digitized analog signals Digital Signal Processors play a central role in the emerging networked world. Used in digital camcorders, answering machines, satellite converters and set-top boxes; digital video disk (DVD) and game players; PCs, modems and hard drives; cellular base stations and wireless phones (Chart 6 & 7); as well as anti-lock brakes, motor control and a host of other emerging applications the DSP market is thriving (**Chart 12**). Texas Instruments leads the market in sales (**Chart 13**) and produces the TMS320C82 which combines two DSPs, a RISC microprocessor with FPU, a memory controller and caches on a single chip device capable of performing more than 1.5 billion operations per second (BOPS) suggesting the potential speed improvements of system-on-a-chip technology.*

The PC will leave home and office and become adult, a global information appliance, to be called the network computer or teleputer. History tells us that in defiance of the fierce determination and creative energies of Microsoft and Intel, this new computer will forge a different computer industry. Such a transition is portentous for the American companies of the old regime who now rule the world and probably will not rule the new one. Andrew Grove in his new book, *Only the Paranoid Survive*, would call it an inflection point.

Grove should know. Seemingly in command of the industry today, with processor speeds rising some 60 percent a year, his regime is now the target of every venture capitalist, microchip entrepreneur, internet projector, and government industrial policy. That's not paranoia; that's a fact of life in the fast lane.

What could be wrong with the paradigm that has made Intel the world's largest chip firm and the US the center of computer progress? With memory dominating the silicon area of every computer, today it makes more sense to put a processor into memory than to put memory on a processor. It also makes more sense to put rudimentary processing instructions into DSP than to put rudimentary DSP on the processor. For a company that has no significant presence in either DSP or DRAM, this reversal is ominous. The fact is that **Texas Instruments**–the world DSP leader and the American DRAM leader–is today better oriented to the new era than Intel is.

All the Pentium functions require access to memory for instructions and data, and memory is still mostly located off chip, reachable only through as few as 16 pins plugged into sockets linked to buses–parallel wire traces inscribed on printed circuit boards eking up in capacity at a rate of only 10 percent a year at best. Finally it all comes down to the pins–the relatively huge and homely copper tabs, numbering in the tens, and each easily visible to the naked eye, that link the microcosm of the chip to the macrocosm of the board. There the infinitesimal miracles of microelectronics halt abruptly and give way to the laborious big world disciplines of wire and solder and plug in sockets.

Intel's response to every problem is to push ahead with faster processors. That is the very nature of a microprocessor company. But there is increasing evidence of diminishing returns for processor speed. Jon Forrest of Berkeley observes that people in general are more satisfied with the speed of their computers than ever before. Overwhelmingly, the problem is not processor speed but input-output. Bob Metcalfe polls his audiences on what they would choose, a faster processor or a faster modem. They are nearly unanimous for a faster modem. They are more interested in saving time on the net than saving time on the desktop. This is a reason why Java's celebrated slowness on routine functions doesn't matter very much; if you are on the net, Java allows you to move around faster than ever before.

Metcalfe's audience responses make sense. Faster Pentiums no longer yield much better performance. The key reasons lie in the defining scarcities of the information age: the speed of light and the span of life.

The speed of light limits the effectiveness of faster processors. Memory bandwidth is growing fast but access times have stagnated. Amdahl's law–that system speed is determined by the slowest components–ordains that latency will rule. As David Clark of MIT puts it, "We can always buy more bandwidth with more money, but we can not buy lower latency. That is determined by the speed of light and you cannot bribe God."

If you look at the future of the computer from the point of view of the CPU, the challenge is to speed up memory. The usual remedy is to create an elaborate memory hierarchy between the fast storage cells on chip and the relatively slow cells in DRAM and use algorithms of probability and locality to move the most likely bits into the fastest memories, called cache. This means using relatively small and expensive static RAMs with lots of pins and complex control logic.

As Greg Papadopoulos of Sun points out, however, if the wanted information is not in the cache just one percent of the time–a phenomenal one percent rate of cache misses–the CPU will typically spend fully half of its time waiting on the memory. Richard Sites of Digital concludes in a *Microprocessor Report* (August 5) survey on the future of the microprocessor that "over the next decade, memory subsystem design will be the *only* important design issue for microprocessors."

The obvious answer to this dilemma is to combine processor and memory on one chip. On chip bandwidth runs in 100s of gigabits per second. With no parallel DRAMs in single in line memory modules (SIMMs), with no memory controllers, bus arbitration, and package pins to manage, delay times can shrink into the picosecond range and power drain can drop by a factor of three.

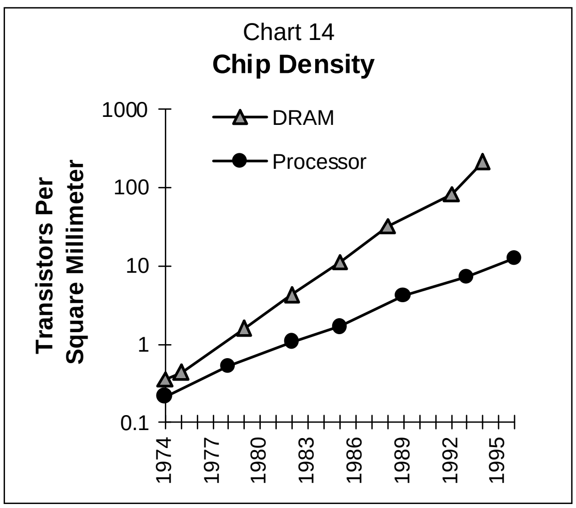

Many companies are already chasing this grail. For example, the ascendant US field programmable logic device producer is **Xilinx**, which combines logic gates with static RAM cells. Intel's processors increasingly incorporate static RAM registers and caches. The appropriate evolutionary path from the point of view of most US chip companies is for the processor gradually to absorb more and more of the memory hierarchy on the chip. Since static RAM is produced by the same essential CMOS fabrication process used to make microprocessors, this strategy means combining processors with static RAM cells.

A typical computer, however, devotes at least 98 percent of its silicon area to memory. The other, less well known of Amdahl's edicts is that for every MIPS of processing power a computer needs roughly one megabyte of memory. A 200 MIPs device with 200 megabytes needs a hundred 16 megabit chips sprawling across the printed circuit board (PCB), occupying at least 50 times as much silicon and board area, pins and wires and power, as the processor itself and choking off communications. The key cost and problem is not the processor; it is the memory.

With simple cells combining one transistor with a tiny capacitor, only DRAMs are dense enough to enable a one chip computer. The essential next step is to move the processor onto the memory rather than the memory onto the processor. In a major coup, **Samsung**'s announcement of an ASIC with a megabit of DRAM cells–moving up to 16 megabits next year– blows away the evolutionary paradigm and brings the teleputer radically nearer. The entire US industry must come to terms with the new reality. The established PC architecture will no longer cut it.

Let us imagine what a teleputer would be like. According to Eric Schmidt, every computer system needs both a brain and lungs. The brain is the microprocessor, but what makes it breathe is the memory system–the lungs. Current computers command superb processing powers but they are gasping for breath through a clogged and constricted windpipe. New network computers do not have to think faster, or think in politically correct ways (i.e. in Wintel instructions and APIs), or summon data from their own memories with dazzling Quiz Kid aplomb. What they have to do is to live and breathe in the ether of the net, with printers, displays, and disk arrays accessed identically whether they are remote or at hand. All have URLs and are called by hyperlinks.

The teleputer should be always on and its browser always up. It sets out from a home page which might be a Netscape Server in Menlo Park, a company server across the country, an **@Home** server or a server at a local ISP. It traverses local and remote storage sites without distinction. Its chief function is to search, find, display, and print information in any form. All its data are objects that bear their own executable codes. The Java programming language began as a vehicle for such appliances. It is a language for computer systems that live and breathe on the net.

Meanwhile, maintenance, upgrades, backups, support will migrate to the ISP, who will serve many of the functions for residential and home office users that the system administrator now performs on office networks. Teleputers could move into schools, hotels, libraries, kiosks, office databases, stores, and remote sales applications. They could be customized by smart cards. They could assume the form factors of digital phones, PDAs, notebooks, news panels, WebTVs, mail readers, and books. Depending upon the expense of the display and the size of memory, they could cost anywhere from a few hundred dollars to the cost of a current workstation. Breaking the model of a unitary computer, however, many of them will not have their

Samsung's announcement of an ASIC with a megabit of DRAM cells blows away the evolutionary paradigm and brings the teleputer radically nearer.

Chart 15

Cable Modem Service Availability

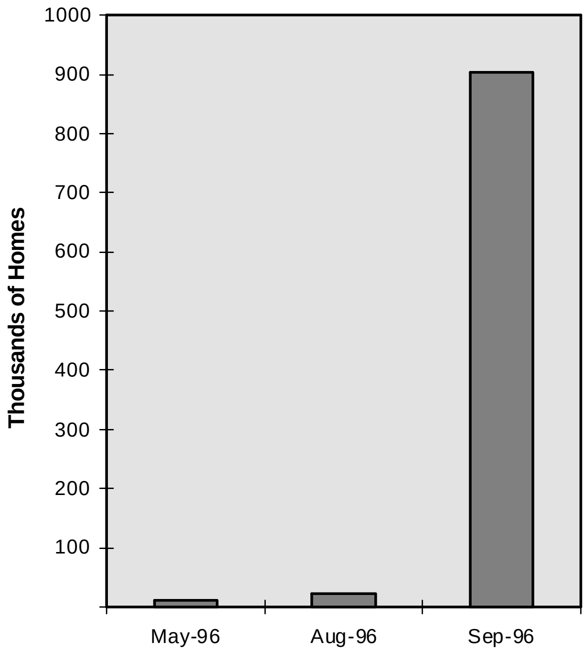

Chart 16

Fiber Deployed by Utility Companies

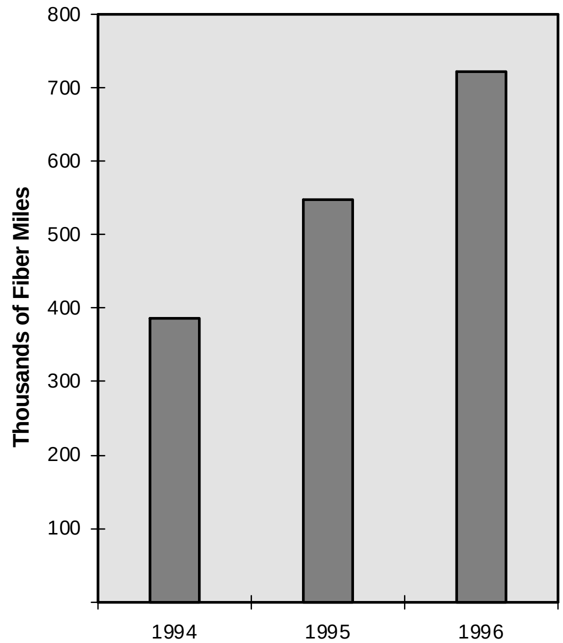

Chart 15. Cable modem service *is now available to over 900,000 homes, with the September start of commercial cable modem services by Time Warner Cable (available to 300,000 homes in Ohio), TCI-@Home (17,250 homes in Fremont, CA), Continental Cablevision (425,000 homes in Jacksonville, Fla.), Rogers Cablesystems (16,000 Newmarket, Ontario, Canada) and several smaller cable operators.*

Chart 16. Utility companies, *owning or controlling the poles, easements and rights-of-way leading to homes and businesses and connecting vast areas of the nation, have the potential to be a major force in telecommunications competing directly with telephone and cable companies. Once concerned only with communications systems to regulate their own power, gas and water delivery grids, utility companies are now broadening their business plans to include data transmission services. Fiber optic cable, immune to much of the electromagnetic interference associated with power transmission, has been the technology of choice for the utilities, and offers them bandwidth to spare. Originally motivated by the cost savings associated with demand-side management to regulate and monitor power usage utilities are increasingly interested in offering a broad range of communications products. Less than two months after the enactment of the Telecommunications Act of 1996, the FCC granted its first application under the Act to allow a public utility to enter telecommunications markets, and has since received 17 additional applications for status as an "exempt telecommunications company."*

own displays. They will invoke available displays through infrared connectors.

Their CPUs will be DSPs, incorporating Java optimized instructions, on large DRAMs interlinked by point to point gigahertz pipes. The network and its standards will prevail in the very heart of the machine. It will solve the problems of the processor memory bottleneck, the maintenance crisis, the complexity slowdown, the ascendancy of realtime tasks, the rise of bandwidth, the centrality of the net. It will be optimized for the horizontal standards of the Internet rather than the vertical standards of the desktop. It will be a structure in tune with the net and in time with the customer and programmer. It is a grown up technology adapted to the new world of nothing but net. It will be a processor that accommodates the lightspeed and lifespan limits that govern the future of information technology.

Amid hundreds of new companies, from **Asymetrix** to **Dimension X**, and established companies, from Borland to **Symantec**, an early beneficiary of the Java revolution will be **Corel**, the currently troubled Ottawa, Canada based company that bought the Word Perfect suite from **Novell** for $153 million after Novell had paid $1.2 billion for it. As a rival to Microsoft in the Wintel arena, Word Perfect had foundered not because the Word Perfect suite is manifestly inferior in functionality to Microsoft Office but because the Windows version was a year late. Corel, however, has already introduced its "Barista" product that allows publishers to create executable content for any Java virtual machine. Now it will be the first to translate its entire portfolio into Java. Microsoft, in all likelihood, will lag more than a year behind. As increasing thousands of corporations move to the Java paradigm with teleputers and intranets, Microsoft may experience the same awful sinking feeling of having missed the train to the future that became familiar at Novell in 1994 and 1995. Of course, as the most prominent Java licensee, Microsoft could give up on its "Captive X" dreams and other proprietary gambits and javatize Office, thus consummating the Java triumph. Checkmate.

Bill Gates long saw this scheme as ridiculous. But with his amazing intellectual openness and strategic decisiveness, he is now moving to make Microsoft the leading Java company. The question is whether he can jettison his proprietary base or whether the proprietary base will corrupt Microsoft's Java products with the usual Redmond gotchas. In any case, the burden lies on those who object to the emergence of what is clearly a needed advance in the industry that would hugely expand the markets for digital technology and software of all compliant varieties. By reducing the maintenance budget by a factor of four or more, it will unleash huge new sales of hardware and software in corporations at a time when many face the crisis of 2000. By offloading crucial costs from home computers, it can ignite huge new sales in homes. In schools, it will benefit from both these gains. The network teleputer is the new computer architecture adapted to the new industrial era.

George Gilder–October 4, 1996

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