GILDER TECHNOLOGY REPORT

February 2000 www.gildertech.com Volume V Number 2

Published Jointly by GILDER TECHNOLOGY GROUP and FORBES MAGAZINE

The Post-Diluvian Paradigm

After the floods of bandwidth, who will greet the dawn and the dove?

What will be the post-diluvian message?

For an analogy that will be familiar to some of you, I will return to the beginning of the twentieth century. At that time, the world of classical physics–all its metrics and models of time and space–collided with the speed of light as an absolute limit. In response, Albert Einstein found himself forced to reconstitute the entire time–space grid of established science.

Today essentially the same thing is happening on the Internet. Over the next five years of the 21st century, the Einsteins of networking will find themselves forced to reconstitute the entire time–space grid of communications infrastructure in similarly elastic four dimensions, bounded on all sides by the speed of light.

Lacking time for the seventeen hops between routers, the long-distance links of the World Wide Web are giving way to all optical systems.

The Einsteinian analogy excludes all minor revisions in the existing system–bars a set of Ptolemaic epicycles or SONET rings or **Cisco** (CSCO) routers upgraded for quality of service. The new Internet will be as radically different from the incumbent as quantum theory was different from Newtonian science.

In This Issue:

Cover Story: The Storewidth Dynamic, Avanex's IPO, Xros' Optical Switch, Mirror Image's Caching Solution

The evidence for the clash with light speed pervades the telecosm list. Lacking time to fetch instructions and data from remote memories, microprocessor boards are moving to single-chip systems (cf., **LSI Logic** (LSI), **National Semiconductor** (NSM), and **Xilinx** (XLXN)). Lacking time for messages to reach remote geosynchronous posts 23,000 miles away, satellites are moving 30 times closer, to low earth orbits (cf., **Globalstar** (GSTRF) and **Loral** (LOR)).

Lacking time for the seventeen hops between routers currently prevalent on the Internet, the long-distance links of the World Wide Web are giving way to all optical systems. Removing most of the delays from the switching and routing process, these ascendant technologies, which have absolutely nothing in common with Cisco routers, will roll out across the cavernous reaches of the Optical Fiber Communication Conference in Baltimore early in March. We will be there to keynote and report. So will **Avanex** (AVNX), our rocketing telecosmic IPO of early February (we all pray that the pop-up $10 billion market cap firm can actually deliver on its promise of dynamically tunable optics before the team cashes out, as is allowed in the company's bizarre agreement with underwriter Morgan-Stanley). Also igniting a blast in Baltimore will be **Xros** (pronounced Kyros), which will dazzle the crowds–get this, *before* announcing an IPO–with a world beating 1,000-port all-optical switch.

Finally, lacking time to travel through all the remote mazes of the net, the World Wide Web will be forced into a new storewidth paradigm imposed by the speed of light limit. Beyond bandwidth abundance, the residual scarcity is latency, and the only way to deal with it is storewidth.

Addressing the fertile interface between bandwidth and storage, storewidth is the accessibility of stored data–the time it takes to find and deliver the first bit of a stored object. I have been touting this paradigm for several months. But most of the companies I have found in the field were well known, such as **IBM** (IBM) or **Novell** (NOVL) (still a storewidth star for its directories), or well valued, such as **Network Appliance** (NTAP) ($19.7B) or **Inktomi** (INKT) ($21B). My original storewidth company, **@Home** (ATHM)–formed to deliver broadband service over cable modems–fell into the maw of **AT&T**

The Internet is a computer on a planet. Like a computer on a motherboard it faces severe problems of memory.

(T) and **TCI**. But under the leadership of technical chief Milo Medin, @Home pioneered the field.

At the heart of the @Home system is ingenious hierarchical memory management and caching to conceal the mazes of slow routers, sluggish switches, and narrowband wires that lurk treacherously within the net. As Medin explains: "The analogy is a single shared memory computer system with multiple processors. You build caches and shared-memory protocols to prevent conflicts. And you mirror and replicate a lot of the data so that it's always available locally. That's what you're going to have to do on the Internet."

In other words, the Internet is a computer on a planet. Like a computer on a motherboard, it faces severe problems of memory access. Thus Internet communications depend on ingenious hierarchical memory management, analogous to a computer's registers, buffers, and latches, its three tiers of speculative caches, its bulk troves of archives, its garbage management systems to filter and weed out redundant or dated data, and its direct-memory access controllers to bypass congested nodes.

In a world of bandwidth abundance, an ever-increasing share of roundtrip delay for a message is attributable to speed of light latency. No matter how capacious the transmission pipes, how large in bits per second the data stream, the first bit in the message cannot move from source to terminal any faster than lightspeed allows, plus the time waiting in queues and buffers at all the switches or other nodes along the way.

Stockholm Syndrome

To most Internet users, the lightspeed limit seems still a secondary issue, and storewidth a triviality. There are many sources of delay more acute, from router conflicts and dropped packets to congested T-1 lines to narrowband Internet Service Providers (ISPs). But the problem of storewidth emerged dramatically in Europe as long as four years ago. As an Internet addict in Stockholm, Sweden, in 1996, Sverker Lindbo was forced to confront the dilemmas of storewidth. As a result, he came up with the essential concepts behind a world-leading storewidth company now called **Mirror Image**.

With Stockholm one of the world's first cities to install a fiber optic network, Lindbo commanded local bandwidth abundance, and Sweden already boasted one of the world's highest rates of Internet use. But just as today, most Internet content was secreted in servers in the U.S. an exorbitant five to eight thousand miles away across the Atlantic.

To bring the Web to Sweden, Lindbo and his colleagues Alexander Vik and Paul Christen developed an entire system of transparent intelligent caching for transferring popular material automatically. When collocated at National Access Points and other Internet hubs around the world, the patented technique, embodied in Content Access Points (CAPs), measures hits, caches pages, and trashes obsolete materials.

With the Net shifting toward interactive and transactional content–originating largely in the U.S. but consumed globally–Mirror Image's Swedish solution is becoming a global imperative. Internet penetration is growing far faster in Europe and Asia than in the U.S. While the U.S. share of the world's Internet users dropped from 70 percent to 54 percent in 1999, the U.S. slightly increased its share of Internet hosts (computers up full time dispensing Internet data) to 80 percent of the global total. Even with no hops or other delays, the light-speed limit alone means that Internet users outside the North American continent are at least 200 milliseconds away from the vast majority of websites.

To fetch a web object using the Internet protocol–whether a frame, image, logo, or banner–takes two to seven round trips between the end user and the Web server. With each page comprising as many as twenty-five objects, those round-trip-speed-of-light milliseconds keep adding up even for entirely static material.

Meanwhile, more and more Internet content is dynamic, consisting of e-business, streaming audio and video, large software and data files, and interactive services such as IP telephony, group games, simulations, and other transactional items. These forms of content often require a round-trip delay of under 150 milliseconds. Uunet guarantees 80 milliseconds to its elite customers. As you can see on Eugene Prescott's website (www.taxtechcpa.com), most of the Internet systems in the world show average delays between 200 milliseconds and 600 milliseconds (Australia), which rapidly multiply from milliseconds to minutes with myriad objects on an obstructed net. As Peter Sevcik has shown in September's *Business Communications Review*, average Internet delay at top business sites has deteriorated since 1995 from 240 milliseconds to 370. The Internet business plans of thousands of dot.com companies cannot succeed unless this deterioration is abruptly overcome. Hence the storewidth paradigm.

Akamai stasis

The speed of light factor imposes four constraints: (1) It prohibits the seventeen hops among routers that the average Internet packet makes before reaching its destination. These seventeen hops consume several times the delay budget for voice and video communications, for example. (2) With a billion Web pages dispersed around the globe and multiplying at a rate of a million a day, comprehensive searches pay a global light delay tax and a complexity tax. The best search engines can cover well under 20 percent of available net contents. (3) Prevailing methods of accelerating Web access do not work for the increasing share of web material that is dynamic and transactional. For example, **Akamai**'s (AKAM) some 2,000 distributed servers concentrate on static signage, banners, frames, tables, titles and the like. (4) Much of the promise of the Web lies in liberating culture from the lowest common denominator programming of television and films. Current systems of distributing these multi-gigabyte files, often in several languages, entail making as many as ten thousand copies and sending them separately around the globe. As long as this broadcast approach prevails, the Web's promise as a TV killer will not be fulfilled.

Bandwidth problems are solved by hardware. If bandwidth is inadequate to handle a particular kind of flow such as streaming video or video teleconferencing, the best solution is to replace the pipes with larger ones and replace the routers with faster ones. The ultimate hardware solution is to create an all-optical path where software is removed entirely and information flashes around the Net on wings of light.

Therefore, as bandwidth increases at the center of the net, software hardens. Driving out the millions of lines of software code in electronic switching and routing systems, all optical networks based on WDM provide a crucial hardware prerequisite of meeting the ever-rising demand for bandwidth. This paradigm is the fibersphere, and it accounts for much of the telecosm list, from **Ciena** (CIEN) and **Nortel** (NT) to **Corning** (GLW) and **JDSUniphase** (JDSU).

As bandwidth and throughput soar, however, speed-of-light latency becomes an ever larger portion of round trip time. While hardware solves bandwidth problems, it can do little to reduce the time for the first bit to be found and fetched. Speed-of-light latency limits must be addressed in software.

While at the center of the Net, software hardens into fiber optic glass, at the edge, hardware must soften. Dumb telephones, TVs, and storage systems based on simple hardwired technology give way to personal computers, Web phones, teleputers, and smart storage that must be customized in software for different forms of data. Software radios are on the way.

Do it with Mirror

The software programs relevant to the light speed crisis deal with searches, directories, caches, and geodesy. The first Internet accelerator that addresses all the key speed-of-light constraints at once is Mirror Image, the company begun by Linbo and Vik in Stockholm in 1996. Currently part of the Internet group called **Xcelera** (XLA), the company plans 32 CAPs collocated in key NAPs and other Internet hubs around the world. Each CAP consists of some $2 million worth of hardware and software, including a Cisco GSR router fronting an **Oracle** (ORCL) 8*i* database that's running, typically, on **Hewlett Packard** (HWP) servers on the CAP edge and core. Mirror Image's proprietary software manages some two terabytes of Internet data that comprises some 75 percent of Web traffic (much of the rest is realtime uncachable information, such as conversations).

Critical to the system is a patented intercept and inject function installed at an ISP or content provider to receive hits and deflect them to the Mirror Image hubs. The first Mirror Image product, now used by some fifty mostly European ISPs, requires a local cache appliance from Network Appliance or Inktomi to divert hits to the Mirror Image CAP. Every time an end user clicks on a web page, it is moved through the ISP's cache to the CAP. The second version now in Beta testing uses a fiber link directly between the Mirror Image router at the CAP and the ISP router's Border Gateway Protocol (WCCP20). Clicks are intercepted and injected into the two terabyte storage system at the CAP, which comprises 75 to 80 percent of all Internet traffic.

Mirror Image customers are ISPs–who pay between $10 and $25 per gigabyte served (depending on volume and geodesic location)–and content providers, who pay $1,000 per month for a peak rate of a megabit per second delivered to their clients. Since all their content is delivered from the CAP, these content customers are saved at least that much in obviated bandwidth costs. So far, the system is up and running in Frankfort; London; St. Louis; Washington, DC; and San Jose. By the end of the year it will be launched in Paris, Amsterdam, Dallas, Chicago, Los Angeles, and New York. And, oh yes–Stockholm.

The first Internet accelerator that addresses all the key speed-of-light constraints at once is Mirror Image.

Why is this system superior to the approaches of the several companies, now worth scores of billions in market cap, that preceded Mirror Image to the fray and that command larger numbers of caches and more complex networking algorithms? The rival accelerators–from Akamai to **Digital Island** (ISLD) to **Adero**–make many of the same promises. The difference is that Mirror Image does it with storewidth, while the others focus on intelligent networking. In an age of bandwidth abundance, intelligent networking is a distraction. By the time all these companies perfect their intelligent routing systems and set up their global private networks, the bandwidth vendors will have moved far beyond them. Total global bandwidth is increasing at a rate approaching tenfold a year. No intelligent routing scheme can begin to keep up. The great insight of Mirror Image is to focus exclusively on storewidth–storing and caching, searching and sorting to deal with light speed latency–rather than on figuring out optimal routes for messages.

For example, Akamai locates a $5,000 **Linux** (LNUX) server at some two thousand ISPs around the globe interconnected by Akamai's network. To store desired Web objects, each of these boxes uses gigabytes of random access silicon memory chips rather than terabytes of disk capacity. Therefore, if you click on an item on the Akamai system, you get it instantly–if it is there. But if the server does not have the desired object, Akamai must go out on the Net and find it. The search function faces all the usual delays and complexities of the net. Using intelligent algorithms to find the fastest routes, the system is good at navigating. But it is ultimately hostage to the speed of light latency and other network problems.

Moreover, Akamai does not readily accommodate the dynamic portions of Web pages. For a relatively small number of objects in great demand, the Akamai system is hugely effective. But the Web

Cash in the Chips

Chart 1

Global IC Market Continues Growth with Record Revenues

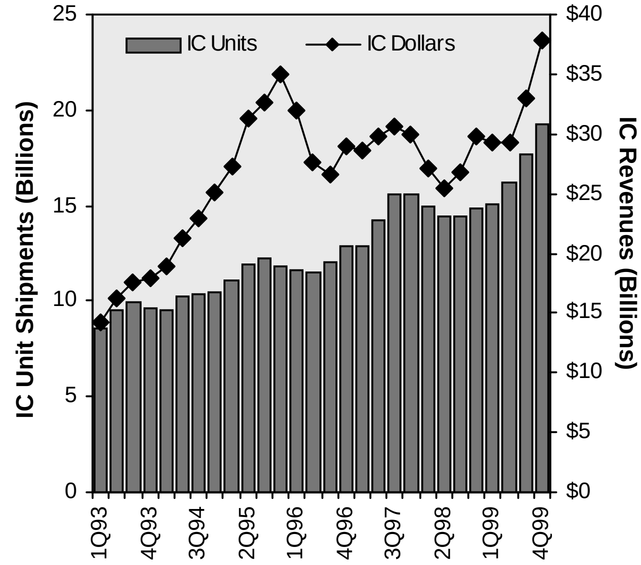

Source: IC Insights

The global semiconductor industry is in the midst of a strong upsurge, with strong growth in units sold and record high revenues (Chart 1). We say it will continue for at least a couple of years. Why? Telecosmic demand feeds Microcosmic prosperity, as communications applications, both Internet and wireless, power chip sales.

To be sure, a strong economy (as the Asian crisis fades) doesn't hurt. But the link between global GDP and IC sales, though powerful, is not consistent. The 1996 semiconductor downturn had more to do with the semiconductor business cycle than the economy. Stable DRAM prices and rising demand led to huge revenue gains, which were reinvested in new fabrication facilities, since most of the world's existing fabs were running full tilt. When a bumper crop of new fabs came online, the capacity shortage became a very expensive surplus, and prices collapsed as manufacturers struggled to support fiercely expensive new fabs. Plunging component prices in early 1996 led us to predict the supply side driven "coming PC boom" which revived the IC industry prior to the economic crisis.

In the wake of the '96 collapse, most industry capital spending went toward improving technology in existing fabs–i.e. more circuits on a chip–rather than new capacity–more chips. (Chart 2). As the PC boom and the Net ignited chip demand, again soaking up worldwide capacity, we predicted the bottom of the cycle in the summer of 1998 and the continued surge we have seen since.

Semiconductor prices are driven in large part by fab capacity utilization (Chart 3). As capacity is constrained average unit (i.e. per chip) selling prices rise and revenues grow; Moore's law helps support an apparently paradoxical insensitivity of demand to rising prices (measured in chip units) by increasing the number of circuits per chip, driving the cost of MIPS (millions of instructions per second) down and the value of chips up.

As always the up cycle is provoking a drive to add capacity and build new fabs, with increases in new capital spending already announced by Intel and others. But construction delays should keep capacity tight, and continued strong demand should keep unit prices high for the next couple of years.

Capacity is strained across the whole range of the industry from the most densely packed cutting edge digital designs to more spacious analog applications. The dark colored portion of the bars in chart 2 represent fab capacity for making chips with feature sizes of greater than 0.7 microns. Producing mostly analog/linear ICs, increasingly for communications applications (Chart 4), capacity utilization in these fabs has surged from 83% to 95%. TI led the analog IC market in 1999 with 11% market share, while National and Analog Devices were in 6th and 7th place with 6% and 5% share respectively.

Chart 2

Recent Capital Spending Focused on Technology, Not New Capacity

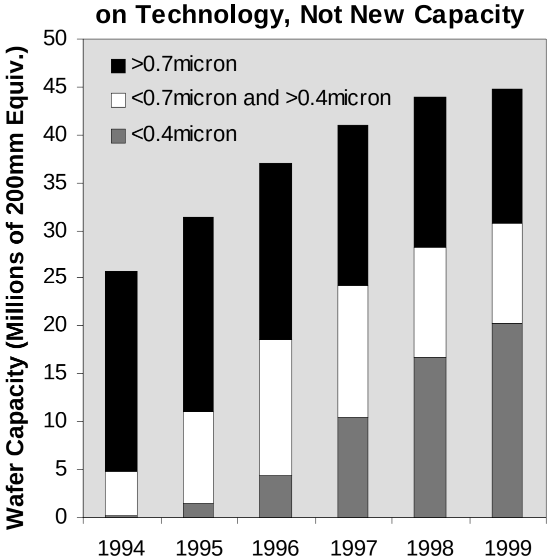

Source: IC Insights

Chart 3

Semiconductor Revenue Growth Linked to Capacity Constraints

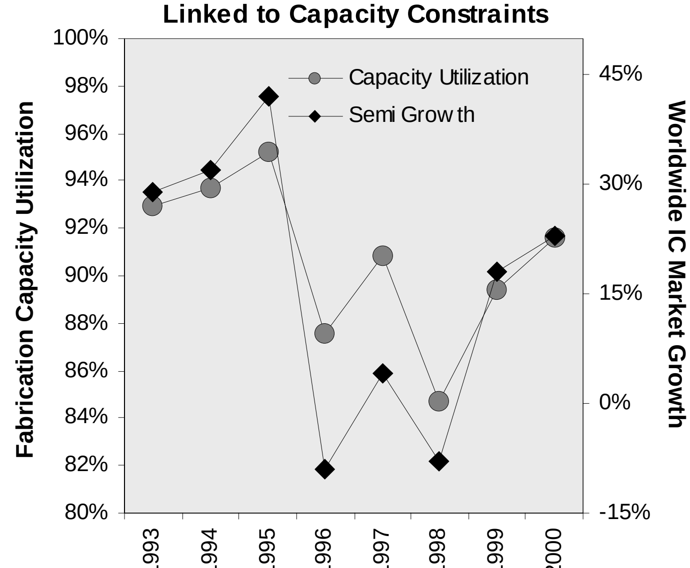

Source: IC Insights

Chart 4

Communications Applications Power Analog IC Market

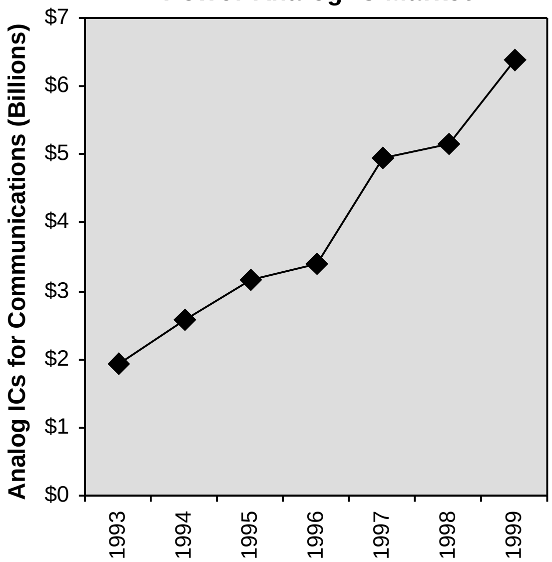

Source: IC Insights

Chart 5

Digital Signal Processor Market Grows with Wireless Applications

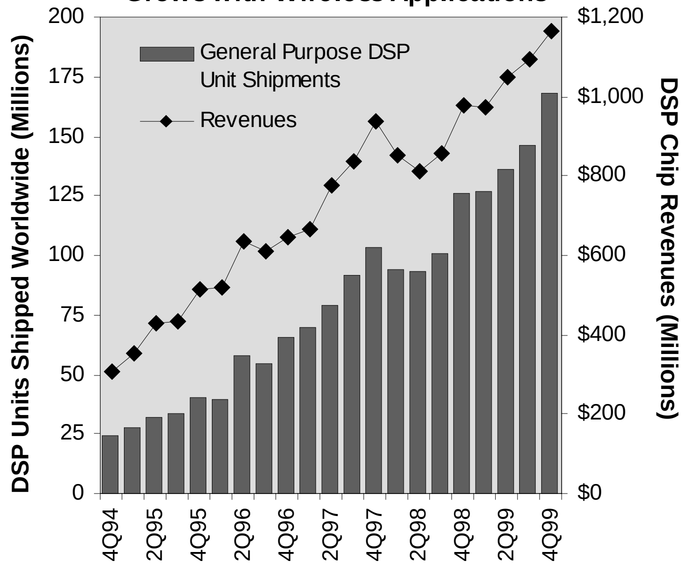

Source: IC Insights

The picture is similar for high end digital chips, with huge densities made feasible by cutting edge sub 0.3 micron processes. Capacity utilization for those fabs has surged from an already daunting 93% to an unheard of 98%. LSI Logic and Lucent were number 1 and 2 in merchant ASIC sales (number 3 IBM would have been first if captive sales are included), while Xilinx and TI were in 5th and 7th place respectively. The demand for wireless ASICs catapulted Qualcomm to the number one spot among fabless IC companies, followed by Xilinx in the number 2 spot, and Broadcom and C-Cube coming in 6th and 8th respectively.

The DSP market continued to soar in both unit shipments and revenues (Chart 5). TI increased its hold on the DSP market increasing its share to 47%, forming what IC Insights describes as an inverted pyramid in which a few players overhang the rest of the market and increasingly dominate specific IC segments (Chart 6). As Bill McClean of IC Insights points out, this model is transforming the IC landscape as barriers to entry rise with the cost of fabs and technology and the value of intellectual property and design expertise. The model is also evident in the programmable logic device market (PLDs, including FPGAs) that is lead by Xilinx (Chart 7). The programmable logic device market highlights the importance of design capability and technology leadership as increased density, complexity, and functionality have caused average selling prices (ASP) and revenues to rise faster than unit shipments (Chart 8).

-Ken Ehrhart

Chart 6

TI Consolidates DSP Lead
1999 DSP Market Share

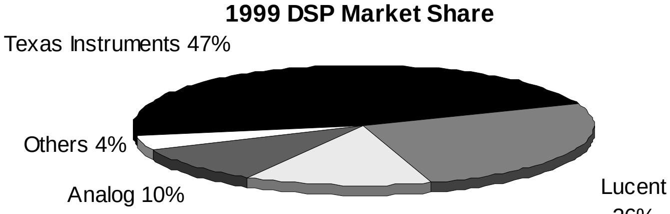

Source: IC Insights

Chart 7

Xilinx Recaptures PLD Lead
1999 MOS PLD Marketshare

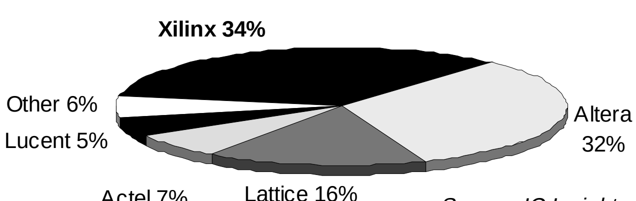

Source: IC Insights

Chart 8

Programmable Logic Device Revenues Outgrow Shipments as ASPs Rise

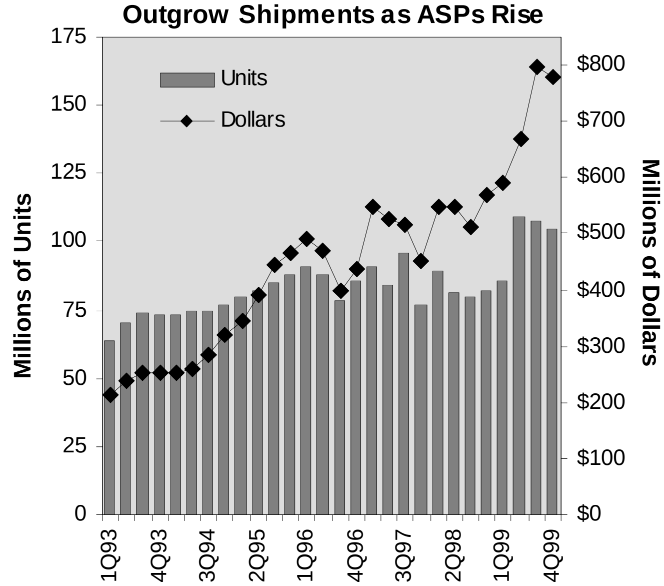

Source: IC Insights

Mirror Image addresses the key problem of bandwidth abundance rather than the residual snags of bandwidth scarcity.

consists of a billion pages mostly in limited demand. Even a popular site under heavy demand, such as a catalog or an encyclopedia, may have many thousands of objects. Akamai can do relatively little for such a customer.

Akamai and its rivals, such as Sandpiper (Digital Island) and Adero are partly in the business of buying bandwidth cheap and selling it dear. Managing links among thousands of distributed caches across the global Internet, these companies must concentrate chiefly on moving data efficiently rather than on storing it accessibly. By contrast, Mirror Image focuses on storewidth at a relatively small number of hubs, keeping its full two terabytes–comprising some twenty million web pages–updated and available at its CAPs around the world. When the system is fully deployed, pages will rarely be more than two or three routers away from the customer. As deployed WDM bandwidth increases, the Mirror Image system will become increasingly effective, for it addresses the key latency problem of a world of bandwidth abundance rather than the residual snags of bandwidth scarcity. Mirror Image is a postdiluvian company.

Distributed around the globe, Mirror Image's thirty-two sites eliminate most of the speed-of-light delay. With the next version of the software, video, audio, and multimedia can be readily supplied with delays well under the obligatory 100 milliseconds. Because 75 percent of web traffic is on each server, searches can be swift and complete compared to the hit-or-miss regime currently in effect at most search engines. The Mirror Image CAPs can handle transactions smoothly because for content customers (now participating in Beta tests) the entire site is mirrored, including all the dynamic algorithms.

The Mirror Image system is vitally needed by Web hosting services, such as **Exodus** (EXDS) or **Global Centers**, that currently provide space for huge amounts of material but do not effectively interlink or process it. These outfits provide real estate and connections for their clients but do not integrate their clients' routers and content with the routers and content of other users of the hub. Thus economies of scale are limited. Mirror Image's customers–ISPs and content providers–are all on one system, so that as the company's business expands, it benefits from Metcalfe's law: the more interconnected users, the greater the share of Internet material will be readily available and the more valuable will be the Mirror Image service. Thus Mirror Image should command a major first mover advantage.

The company is currently owned by Xcelera, which commands a market cap of some $3.6 billion. It owns a group of interesting Web companies, mostly oriented toward Europe, which is the fastest growing Internet region. Included in Xcelera are Deo.com, a music on-line service; MNW records, a music content provider; e-game, an ad-supported game company used in Europe by **Microsoft** (MSFT), **Procter and Gamble** (PG), and **Coke** (COKE), among others; OneSure.com, an online insurance provider; Wideyes.com, an Internet recruiting service; and CoreChange, a Java-based portal software program that puts corporations on the net with their own portals.

CEO Vik, however, knows that Mirror Image, headquartered in Woburn, Massachusetts, is his chance to change the world. He plans to spin it off in an IPO that can finance the creation of the full thirty-two CAPs. Conveying the full power of the storewidth paradigm, Mirror Image joins our list this month.

Terayon Over the Rainbow

Every GTR subscriber knows a company that developed CDMA technology that others first claimed was unnecessary and unworkable. The technology proved itself with better handling of data and greater capacity. It excelled in Asia, and was finally adopted as an industry standard. The name of that company, of course, is…**Terayon** (TERN), now poised to repeat **Qualcomm**'s (QCOM) remarkable success.

Terayon was formed in 1993 in Silicon Valley to pursue the vision of broadband over cable developed by Zaki and Shlomo Rakib in Israel. Analogous to Qualcomm's wireless CDMA, S-CDMA uses spread spectrum techniques to modulate signals across a swath of available spectrum rather than chopping them up into time slots in a narrow band as in Time Division Multiple Access (TDMA). Just as in wireless, S-CDMA over cable offers a shared medium system with more flexible handling of bursty data, greater capacity, and far superior immunity to noise. Like urban cellular environments, cable coax spectrum is fraught with interference and noise.

Terayon, however, clashed repeatedly with the DOCSIS group setting cable modem standards. DOCSIS wanted the company to surrender its technology into a royalty-free pool of patents to be shared by any company developing products for the standard. Zaki refused to surrender tiny Terayon's key asset. Just as the European and U.S. wireless standards were adopted well ahead of Qualcomm's participation, the cable modem standard emerged without Terayon.

By 1998 Terayon was nearly broke, when Canadian cable company **Shaw Communications** (SJR)–which saw S-CDMA as a critical edge in converting cable TV customers to broadband Internet users–approached the Rakibs to make an investment. With Terayon's technology Shaw leads all @Home partners in subscribers and cable modem penetration rates. Shaw documented Terayon's unique immunity to interference from outside sources ("ingress" is the unfortunate industry term) a crucial source of noise and thus degraded signal-to-noise ratios (SNR). TDMA modems will function only in robust cable systems expensively upgraded to maintain a certified signal-to-noise ratio of 35dB (some 7,000 to 1) by adding noise filters, replacing a greater portion of coax with fiber optic links, and reducing the number of homes (and compounded noise) per node. S-CDMA modems, requiring less than 15 dB of SNR at full capacity, can thus run on a cable plant certified to only 25dB, like Shaw's, and still maintain a 10 dB operating margin. Because of S-CDMA's superior noise handling, Shaw was able to avoid upgrade expenses and speed rollout to nodes servicing as many as 12,000 homes (an order of magnitude more than rivals).

Shaw's success caught the attention of the other great Canadian cable provider, **Rogers** (RG), which had been using Nortel cable modems from LAN City (Bay). Rogers turned to Terayon for help. Terayon struck two deals: one with Rogers cable for Terayon's cable modem system, and one with Rogers Communications for a joint venture to develop voice over cable technology. Almost immediately, according to Zaki, the Rogers voice partnership confirmed that VoIP for cable was not ready, particularly with TDMA. But Terayon acquired an Israeli cable telephony firm called Telegate that was already using Shlomo's S-CDMA.

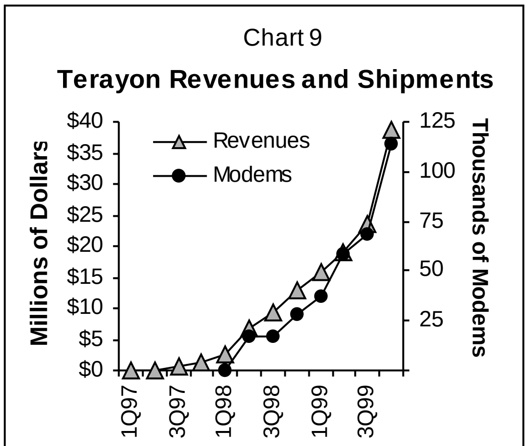

Source: Company Announcements

The Terayon acquisition strategy follows a two dimensional grid, with one axis for content (including data, voice, and video) and the other for conduit (including cable, DSL, and wireless). The first square–data over cable–was filled by Terayon's original cable modem technology. The voice-over-cable square was filled with the Telegate purchase. And video over cable was achieved with the Imedia buy. The recent addition of DSL technology from Radwiz and **ANE** (Raychem Access Network Electronics) moves all squares one column to the right into the DSL realm, where Terayon will begin offering, within weeks, an integrated product combining multiple voice lines with always-on data. Imedia's video multiplexing will follow, completing the data, voice, video column for DSL and allowing Terayon to challenge **Next Level Communications** (NXTV), the current leader in DSL-based video. Purchase of privately held **Combox** of Israel moved Terayon into wireless through Combox's strength in DVB (digital video broadcasting). This is an emerging satellite broadcast standard adopted by Loral and others that is also the basis for a European cable modem standard that is competing with the U.S. DOCSIS standard and Terayon's proprietary S-CDMA.

Terayon's Imedia CherryPicker is a digital video multiplexer of amazing flexibility that can mix and match satellite network feeds, local content, and ads. Desired channel transmissions can be scaled according to content, from a high-definition Superbowl signal with a visibly spiraling football, to a narrowband offering of newscasters motionlessly reading the latest headlines. Cable companies can groom the broadcast stream and insert local news bulletins and regional advertisements within nationally syndicated programming. Already a hit in multilingual regions like Switzerland for supplying programs with selectable language tracks, the Imedia system will ultimately scale downward. It can target ads to the diverse desires, demographics, needs, and interests of individuals or groom DSL's limited video bandwidth according to the needs of the viewer.

Meanwhile, Terayon's semiconductor business developing and producing S-CDMA modem chips is thriving. Along with **Broadcom** (BRCM), **TI**, and **Conexant** (CNXT), Terayon is among the few companies producing cable modem ICs. S-CDMA is targeted for incorporation as in advanced physcial layer in the U.S. cable modem standard, DOCSIS.

Already number one in Canada and Japan, Terayon now hopes to expand its dominance into the rest of Asia. Ken Ehrhart was in Hong Kong in January as Terayon and Hong Kong's **I-Cable Communications** (ICAB) announced a broadband Internet partnership. Vince Lam, I-Cable's technical chief, chose Terayon modems after several years of extensive tests with systems from IBM, Nortel's LANCity and Com 21, as well as Terayon, showed S-CDMA's decisive superiority. With cable modem Internet access under way, I-Cable plans to offer voice-over-cable services to over 90 percent of Hong Kong residences. One vision is for I-Cable to act as an @Home of Asia, with Terayon providing the equipment. With deals or deployments in Hong Kong, China, Japan, and Korea, Zaki Rakib envisions the possibility of developing an Asian cable modem standard based on S-CDMA's superiority in high-density urban settings.

I-Cable's charismatic CEO, Stephen Ng, remarked at the press conference that, several years ago, Wharf joined with **SingTel** and Qualcomm to present a joint bid based on CDMA technology for wireless licenses in Hong Kong, but the government denied the bid. Now he is placing a new bet–on Terayon's S-CDMA–and is looking to make it the standard throughout China and other parts of Asia. Terayon's fourth quarter results showed shipments and revenues up over 65 percent over the previous quarter, with a proforma profit of four cents per share (compared with the analysts' consensus of a nineteen cent loss). If you failed to place your bets on Qualcomm, you, too, have another chance for a CDMA star, Terayon.

–George Gilder (with Ken Ehrhart)
February 15, 2000

If you failed to place your bets on Qualcomm, you have another chance for a CDMA star, Terayon.

TELECOSM TECHNOLOGIES

ASCENDANT TECHNOLOGY	COMPANY (SYMBOL)	REFERENCE DATE	REFERENCE PRICE	JAN-'00: MONTH END	52 WEEK RANGE	MARKET CAP.
CABLE TECHNOLOGIES/SERVICES						
Cable Modem Chipsets	Broadcom Corporation (BRCM)	4/17/98	12 *	289 5/16	46 1/4 - 325 5/8	30.117B
S-CDMA Cable Modems	Terayon (TERN)	12/3/98	31 5/8	107	25 3/4 - 129 1/2	2.333B
MICROCHIP TECHNOLOGIES						
Analog, Digital, and Mixed Signal Processors	Analog Devices (ADI)	7/31/97	22 3/8	93 1/2	24 3/8 - 104 1/2	16.409B
Silicon Germanium (SiGe) based photonic devices	Applied Micro Circuits (AMCC)	7/31/98	11 11/32	147 3/4	16 - 151 1/4	8.008B
Programmable Logic, SiGe, Single-Chip Systems	Atmel (ATML)	4/3/98	8 27/32	31 1/16	6 13/16 - 35	6.262B
Digital Video Codecs	C-Cube (CUBE)	4/25/97	23	69 3/16	17 1/4 - 70 1/4	2.835B
Linear CDMA Power Amplifiers, Cable Modems	Conexant (CNXT)	3/31/99	13 27/32	84 1/2	7 1/2 - 87 3/8	16.469B
Single Chip ASIC Systems, CDMA Chip Sets	LSI Logic (LSI)	7/31/97	31 1/2	81 1/2	21 - 83 5/16	12.119B
Single-Chip Systems, Silicon Germanium (SiGe) Chips	National Semiconductor (NSM)	7/31/97	31 1/2	52 1/2	8 7/8 - 52 1/2	9.083B
Analog, Digital, and Mixed Signal Processors, Micromirrors	Texas Instruments (TXN)	11/7/96	11 7/8	107 3/4	43 - 120 5/8	85.381B
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	10/25/96	8 7/32	45 3/4	16 1/4 - 46	14.699B
OPTICAL NETWORKING						
Wave Division Multiplexing (WDM) Systems, Components	Ciena (CIEN)	10/9/98	8 9/16	65 5/8	16 5/8 - 72 1/4	9.069B
Optical Fiber, Photonic Components	Corning (GLW)	5/1/98	40 15/16	154 1/4	43 9/16 - 156 1/4	39.766B
Submarine Fiber Optic Networks	Global Crossing (GBLX)	10/30/98	14 13/16	50 3/4	20 1/4 - 64 1/4	40.336B
Wave Division Multiplexing (WDM) Components	JDS Uniphase (JDSU)	6/27/97	7 1/4	203 15/16	18 - 248 1/2	58.490B
Broadband Fiber Network	Level 3 (LVLT)	4/3/98	31 1/4	117 5/16	45 1/4 - 120	40.229B
Wireless, Fiber Optic Telecom Chips, Equipment, Systems	Lucent Technologies (LU)	11/7/96	11 25/32	55 1/2	47 - 84 3/16	174.270B
Broadband Fiber Network	Metromedia FIber Network (MFNX)	9/30/99	24 1/2	67 11/16	17 3/8 - 72 3/8	15.756B
Wireless, Fiber Optic, Cable Equipment, Systems	Nortel Networks (NT)	11/3/97	23	95 1/2	26 15/16 - 110	129.880B
Broadband Fiber Network	NorthEast Optic Network (NOPT)	6/30/99	15 1/16	98	10 1/2 - 103	1.597B
WIRELESS TECHNOLOGIES/SERVICES						
Low Earth Orbit Satellite (LEOS) Wireless Transmission	Globalstar (GSTRF)	8/29/96	11 7/8	32 5/16	12 5/8 - 53 3/4	3.156B
Satellite Technology	Loral (LOR)	7/30/99	18 7/8	19 5/8	13 1/2 - 25 3/4	4.806B
Nationwide Fiber and Broadband Wireless Networks	Nextlink (NXLK)	2/11/99	20 7/16	84 3/8	17 7/8 - 101 1/8	11.230B
Code Division Multiple Access (CDMA) Chips, Phones	Qualcomm (QCOM)	7/19/96	4 3/4	127	7 1/4 - 200	89.941B
Nationwide CDMA Wireless Network	Sprint PCS (PCS)	12/3/98	7 3/16	55 1/32	28 - 114 3/8	51.139B
Broadband Wireless Services	Teligent (TGNT)	11/21/97	21 1/2 *	67 1/16	33 3/8 - 90 3/4	3.628B
INTERNET TECHNOLOGIES/SERVICES						
Internet Enabled Business Management Software, Java	Intentia (Stockholm Exchange)	4/3/98	29	20 3/4	17 1/2 - 35 1/4	0.496B
Network storage and caching solutions	Mirror Image (Xcelera) (XLA)	1/31/00	116	116	5/8 - 158 7/8	3.062B
Telecommunication Networks, Internet Access	MCI WorldCom (WCOM)	8/29/97	19 61/64	45 15/16	40 5/8 - 64 1/2	130.470B
Directory, Network Storage	Novell (NOVL)	11/30/99	19 1/2	33 3/8	16 1/16 - 42 7/16	10.897B
Java Programming Language, Internet Servers	Sun Microsystems (SUNW)	8/13/96	13 3/4	78 9/16	23 1/8 - 84 1/2	137.480B

* INITIAL PUBLIC OFFERING

NOTE: This table lists technologies in the Gilder Paradigm, and representative companies that possess the ascendant technologies. But by no means are the technologies exclusive to these companies. In keeping with our objective of providing a technology strategy report, companies appear on this list only for these core competencies, without any judgement of market price or timing. Reference Price is a company's closing stock price on the Reference Date, the date on which the company was added to the Table. Since March 1999, all "current" stock prices and new Reference Prices/Dates are closing prices for the last trading day of the month prior to publication. Mr. Gilder and other GTR staff may hold positions in some or all stocks listed.

GILDER TECHNOLOGY REPORT
published by Gilder Technology Group, Inc. and Forbes Inc.
Monument Mills • PO Box 660 • Housatonic, MA 01236
Tel: (413)274-3050 • Fax: (413)274-3031 • Email: gtg@gildertech.com
EDITOR: George Gilder
PUBLISHER: Richard Vigilante
DIRECTOR OF RESEARCH: Kenneth Ehrhart
TECHNOLOGY ANALYST: Jeff Dahlberg
MANAGING EDITOR: David S. Dortman
SR. MARKETING MANAGER: Rosaline Fernandes
BUSINESS DEVELOPMENT: Mark T. Ziebarth

FOR SUBSCRIPTION INFORMATION TELEPHONE TOLL FREE:
(888) 647-7304
WEBSITE:
www.gildertech.com