GILDER TECHNOLOGY REPORT
Published jointly by Gilder Publishing, LLC and Forbes Magazine www.gildertech.com February 2001/Vol. VI No. 2

The Storewidth Warp

BlueArc imperils all the software based network storage appliances, whether it be from NetApp, EMC, HP, Procom or Compaq

Exhausted by the travails of time and space, Charlie Burger's jet-lagged body craved the San Jose Hyatt. Burger as you know is our main man on optics, and a key fact of optics is that photons cannot be stored. They are pure moving energy and by nature cannot be settled in a capacitor or trapped in an insulated gate like an electronic charge in a silicon memory, or durably defined like a magnetized domain on a disk, or deposited comfortably at a Hyatt or a Motel 6, even if management leaves a light on for them. Although *Science* magazine recently recounted experiments in which photons were slowed down and even stopped for a millisecond or so in a laboratory, a close reading revealed that in the process the photons died.

Charlie is a lot like a photon. But he also has some 155 pounds of mass, honed to a hard heft by a regimen of slow power lifting. After 18 hours of airports and rental cars and interviews on optics he needed to stop for somewhat more than a millisecond and consign himself to a good night's storage, even if he had to undergo a discrete cosine transform and curl up compressed on a 19 inch rack in an **Exodus** (EXDS) data center.

But it was not to happen. Before Charlie's head could hit a pillow, he would be abducted by aliens and taken on a trip into the storewidth warp, where he would discover what may prove to be the most disruptive company in the history of this letter. Within a week, as a freak effect of global warming, snow would fall in Silicon Valley and the entire storewidth landscape would be frozen in anticipation of a **BlueArc** from outer space.

Living the time-space dilemmas of storewidth, Charlie was torn between optics and storage. A time-bound activity, storage enables movement of information from one moment for use in another. The bits go somewhere and stay put until they are consumed serially by time-bound humans. Communication is space travel. It moves information through space, from one place to another.

With many transmissions possible between the same two points at once, communication ideally is parallel and real time. The ultimate parallelism is wavelength division multiplexing (WDM)—putting data beams of many colors at once on a single fiber thread and moving them across the land at lightspeed. Between time and space is the storewidth warp.

In the scheme of abundances and scarcities that shapes the economics of technology, all computing can be mapped to this matrix of time and space. The usual computer architecture

In This Issue:

embodied in the typical Pentium is a time machine. With over 90 percent of the chip devoted to various forms of memory, the actual computation uses a tiny portion of the chip. But it uses this tiny space repeatedly at a frenzied pace to execute a series of software instructions on a flow of data words fetched from memory. It cuts time into nanosecond slices—billions of cycles every second—and if **Intel**'s (INTC) heroic chip designers have their way, two gigahertz Pentiums will soon slice time into a few hundred picoseconds (trillionths of seconds).

Emblazoned in this architecture is the idea that space is scarce. By spreading out computations over time, serial processing economizes on space. All calculations are done in one place, the central processing unit (CPU), one after another in time. Since only one calculation can occur at any moment in time, speed is gained by shortening the moments.

Beset by ever lengthening TDM user queues on the network side, web servers face the demands of heterogeneous data on the storage side

Parallel processing, by contrast, economizes on time by spreading computations out over space. Rather than conducting the operations in the CPU, the CPU stands to the side and controls the pattern of links through which the data flows. A data flow machine resembles a piano, with the piano wires representing the logic gates and the keys representing the control inputs. As Carver Mead and Lynn Conway put it in *Introduction to VLSI*, "A complex function may thus be performed...by the data path just as the static-appearing array of piano wires may produce a complex and abstract piece of music when a series of notes and chords are struck in a particular order."

Tredennick's paradigm

The fibersphere is such a space machine. Wavelength division multiplexing is inherently parallel and pianistic. You can't store photons. As wavelengths multiply, they fly in parallel through the fibersphere. By proliferating lambdas–optical pathways or circuits–photonic networks enable many users to share the same fiber pathways yet still have their own exclusive circuits. In the fibersphere there is no time sharing or time division, no TDM process of distributing the data packets of many users in time slots on the circuit.

As the fibersphere explodes into ever more parallel paths, giving more and more users access to circuits, computers at the edge can no longer keep up by slicing time into smaller moments. Even as massively parallel optical bandwidth overwhelms the TDM and SONET electronic switches in the core of the network, massively parallel throngs of users jam the serial electronic servers on the edge. Already beset by ever lengthening TDM user queues on the network side, web servers face the increasingly complex demands of heterogenous data on the storage side.

Crucial to overcoming the bottleneck between storage and fibersphere–between time-bound electronics and spacious optics, between you and the Web–is the development of space machines on the edge. Needed are storewidth servers that can match the parallel flood of photons with parallel paths of electronic processing. Relinquished must be some of the general purpose flexibility of the microprocessor. But today the microprocessor era is ending. All technologies must adapt to the awesome power of the fibersphere.

Avanex accessory to kidnapping

Before Charlie encountered his storewidth epiphany, he gained new appreciation for the power of this optical technology that was rapidly transforming the Peninsula into a Silica Valley of glass. Capping Charlie's day, Simon Cao of **Avanex** (AVNX) revealed his plans for a spectacular presentation of the Metro PowerExpress and PowerExchanger at the March 2001 Optical Fiber Conference (OFC) in Anaheim. The seamless metropolitan area network would cut the fiber ribbon, so Simon implied, on an era of networks with thousands of parallel lambdas spread across the Web without any switches in sight.

How could networks do without switches? Charlie envisaged long lines of **Cisco** (CSCO) and **Nortel** (NT) engineers jumping off the San Mateo Bridge. Networks almost by definition had to be switched. But switches imply buffering and storage along the path and thus intolerable delays on the network itself. Simon's wavelength division multiplexed networks would ride lambdas (wavelengths) that could be added or dropped optically, muxed and demuxed, without an actual cross connection. At the beginning of the trip you choose your lambda and the light does the rest. If you had fiber everywhere, tunable lasers, filters, couplers, and splitters, and a lambda for every route, it could work. If I had any bread, I could make a ham sandwich, if I had any ham.

Undaunted by the admittedly long run challenge of removing switches altogether, Cao is launching his concept in the metropolitan area network where it is already taking hold. The campaign could take a decade. Perhaps a switch or two could be sold in the interim. (The optical switch guys could come in off the bridge).

In any case, absorbing the implications had overloaded Charlie's processor for the day, and he had begun his descent somewhere over the WDM rainbow. In the parking lot his Dollar Rent-a-Car reflected the long seven hundred nanometer rays of the California dusk. The day's final hard drive beckoned.

Easy to please, Charlie would have settled for simple and stupid highway attached storage, without queues and buffers in the lobby or tricky transactional protocols at the front desk. He was ready to go without guaranteed quality of service or redundancy (one bed would do) and without chocolate wafers on his pillow, data masseuses on his disks, or 50 millisecond restoration if the California lights browned out.

But what's this? As he walked toward his car, he became aware of a distraction. "No," demurred Avanex's VP of investor relations, Tony Florence, "over here." He pointed to a limo with a chauffeur in a black suit and dark sunglasses. Gently ushered into the black Lincoln Towncar, Charlie heard Tony's last words as the door shut: "Ask for Pesatori, code word *Alessandrini*."

Straining to see the dim lights through the tinted windows, Charlie asked his driver where they were headed.

"North, I believe."

At last Charlie caught a glimpse of a sign. "South" was all he could make out. Forty minutes later he found himself in the parking lot of what appeared to be a deserted office building. Definitely not the Hyatt. He knocked at the entrance.

From behind him came a sharp, "I'll wait here. You may need me." Generous, he thought, of his dour driver. In the darkened foyer, one of what seemed to Charlie to be Earth's last two employees opened the locked door. Charlie tried to summon the code. "Pesadrini, Allesantori…. Synaxia… no BlueArc… but don't tell anyone…. Whatever." His codewords had no effect. Charlie turned on his heels to leave.

"Stop!"

Charlie froze mid-step.

"Down the hall. They've been expecting you." He followed his guide through the empty corridor into a dimly-lit office. On the right, three men encircled an imposing table. "Enrico Pesatori—friend of Walter Alessandrini," one announced himself in a deep, old country voice. "We're in stealth mode," he said, superfluously. Charlie glanced back for his chauffeur. The man was gone. The door clicked, and a revolution in storewidth unfolded.

Silicon storewidth

Whether under its original name Synaxia, or the consultant's new moniker BlueArc, the storage devices on the edge of the network would at last conform to the naturally parallel optics in the center. Matching the parallelism of the fibersphere would be a parallel serversphere. In the storewidth warp, time and space would converge. The Internet would never be the same again.

Hey, you've heard that before, when Kim Polese first shook her Castanet at Telecosm, when Amazon began giving away "Champagne smoochies soap and aromatic votives" on its website, when Pointcast launched the vaunted age of "Push" in *Wired*, and when Matt Drudge levered the political landscape from a trashcan on the Web. You heard it too when digital subscriber line (DSL) and cable modem technologies were unleashed to bring broadband access to the Web for all, guaranteed by promises from politicians and commissioners in Washington and the fifty states. You heard it before. Now it is truly happening–and none too soon…

Even with a broadband local loop in place, the Net will encounter a bottleneck in the persistence of old, time-bound computer architectures in the face of the space travel of optics. The result is a jam at the storage file servers that deliver web pages and other programming to the Net. A web server is a robust general purpose computer that can operate 24 hours a day on any of thousands of different software programs devoted to communications protocols stacks, authentication and security rules, file transfer protocols, messaging and transactions schemes, all retrieved from locally attached disks.

Until recently servers were actually faster than the network, so you didn't have to worry about them. But with the arrival of WDM in the backbone and gigabit and multi-gigabit Ethernet optics in the metro, campus, and LAN, storage file server speeds have become the Web's major sticking point.

> **The faster Ethernet became, the more of a bottleneck the NAS device in its path turned out to be**

The bottleneck comes less from the processor itself, which has moved over the last decade from scores of megahertz to the gigahertz, than from the internal buses, which move data bits around the computer backplane for processing. Since 1993, Pentiums have connected to memory through the PCI (Peripheral Component Interconnect) bus, which has increased its speed merely from 100 megahertz to 133 megahertz during that period.

Bandwidth inside the computer has fallen behind bandwidth outside the computer; outside optical bandwidth is growing some ten times faster. It was superior input-output (I/O) speeds that put the computer in a box with its storage in the first place. As optical network speeds destroy the same box advantage, the computer is disaggregating across the Net. Since a computer is essentially an I/O device for accessing and processing storage, the movement of storage to the Net hollows out the computer and makes way for BlueArc.

The BlueArc revolution

It turned out that BlueArc had nabbed the wrong *GTR* analyst. Charlie arrived back in Housatonic, eyes glazed over, mumbling something about tinted windows and sequential Iometers, insisting we call a gentleman named Enrico. Determined to get to the bottom of things, we put in a call to this Enrico, and set up a meeting.

One week later, in New York at the Gerson Lehrman Group offices in Manhattan, which provide crucial research to the GTR through their huge database of technical experts, we finally met Charlie's storewidth angels. Present was not only CEO Enrico Pesatori, but also Dr. Geoff Barrall, CTO, of the BlueArc Corporation. Neither Charlie nor I had any clear idea what a storewidth bonanza he had uncovered during his magic carpet ride in Silicon Valley.

Two years ago, working to improve Ethernet switches, Dr. Barrall, a Ph.D. in Cybernetics, and his colleagues in a U.K.-based consultancy quickly discovered that success was fruitless. The faster Ethernet became, the more of a bottleneck a network attached storage (NAS) device

DISK PERFORMANCE DEGRADATION

Exponentially mounting data supplies have driven the demand for storage capacity through the roof. Historically, disk manufacturers have heroically met the challenge of providing consumers with smaller, faster, and cheaper forms of disk storage.

In the past five years, magnetic disk storage has experienced its steepest price decline in forty-three years (Chart 1), making magnetic disk storage the data storage method of choice (Chart 2).

In the early 1990's disk read/write heads soared more than 90 nanometers (billionths of a meter) above the disk's surface (Chart 3). Head technology improvements over the past decade have brought the physical distance between the head and the disk surface down below 10 nanometers. Advanced head technologies have led to bit cell size and disk trackwidth reductions (Chart 4).

Smaller, faster, cheaper inescapably leads to denser, smaller bits on narrower tracks. Decreasing the size of a bit cell by a factor of two quadruples disk density. Areal density, the product of track density and the number of bits per track, has increased in by a factor of three million in just over four decades (Chart 5).

Disk performance, however, has not maintained the Moore's Law-pace of disk densities. Densely packed bits mean fewer disks for a given amount of storage and thus fewer drive heads per gigabyte. Time to access a gigabyte, therefore, is actually deteriorating as seek times improve at a relatively slow linear rate while densities surge exponentially (Chart 6).

Disks have become a performance bottleneck. High performance alternatives to magnetic disk storage have a strong growth potential. Solid-state disk storage from companies like Solid Data and Quantum (HDD) addresses the disk drive I/O bottleneck head-on by caching small highly transactional data in a separate cache to improve overall disk drive performance. InPhase Technologies, launched last week out of Lucent's Bell Labs, will address the magnetic disk bottleneck by focusing on holographic storage.

-Mary Collins

Disk storage prices have declined by a factor of 30,000 in the past 40 years.

Chart 1

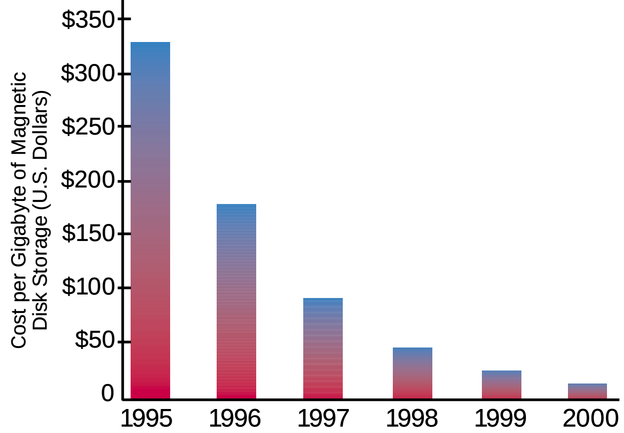

Data supplies have driven storage sales through the roof.

Chart 2

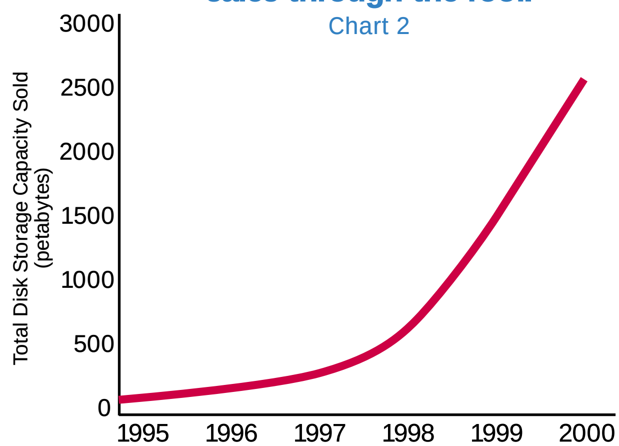

Data Sources: IBM, International Data Corporation

---

placed in its path turned out to be. Network attached file servers had not yet entered the optical age.

High-speed Ethernet demanded high-speed storage servers. Yet all existing NAS appliance designs were based on traditional server architectures. Designed for computation, they become I/O bound in today's high throughput networks. Two gigabit Ethernet (2,000 Mbps) coming in off the LAN and 1.7 gigabit per second fibre channel at the storage end swamp even the fastest storage servers on the market, such as the NetApp 840 and the EMC IP-4700. These software intensive devices traditionally run at 200 Mbps to 300 Mbps, 400 Mbps when streaming audio or video.

Under Amdahl's law, high speed server components can operate no faster than the slowest link on the data path. The whole is *not* better than the sum of the parts. A one-gigabit per second PCI bus (I/O interface), for example, can slow to a standstill while main memory and CPU handle one customer at a time. They step and fetch software instructions for the Internet's TCP/IP (Transport Control Protocol/Internet Protocol) stack, then shuttle back and forth from memory to carry out Windows and Unix file transfer protocols. Upgrading your NAS from a one thousand dollar system to a one hundred thousand dollar system buys a lot more storage capacity but affords essentially the same throughput.

Barrall got together with Pesatori, a former Digital Equipment and **Compaq** (CPQ) VP and president of Tandem Computers (now part of Compaq). Together they assembled a team of sixty engineers to build a multi-gigabit throughput server from the ground up.

A server capable of reading and writing simultaneously at

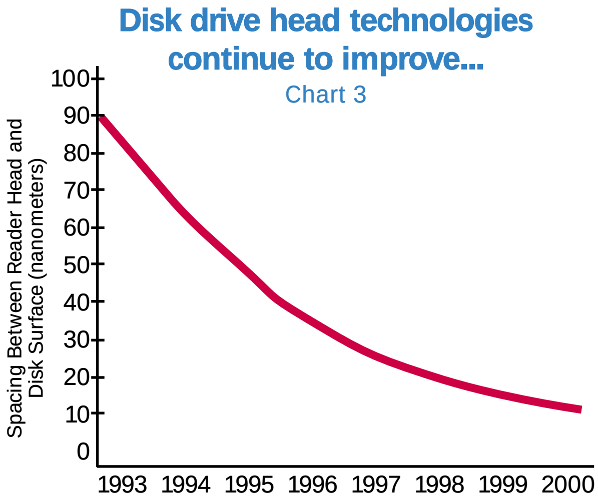

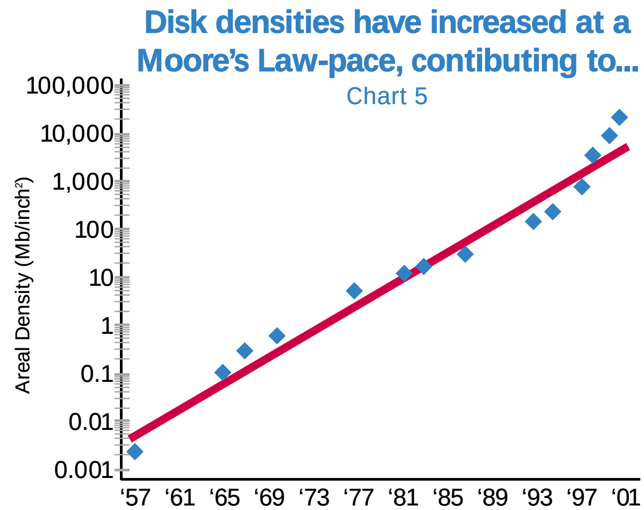

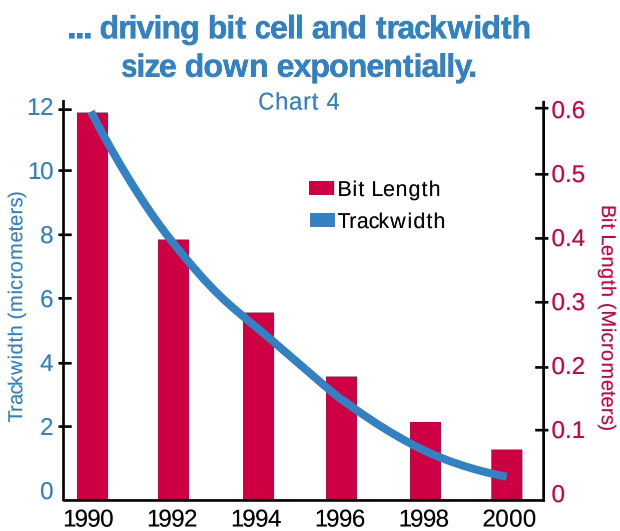

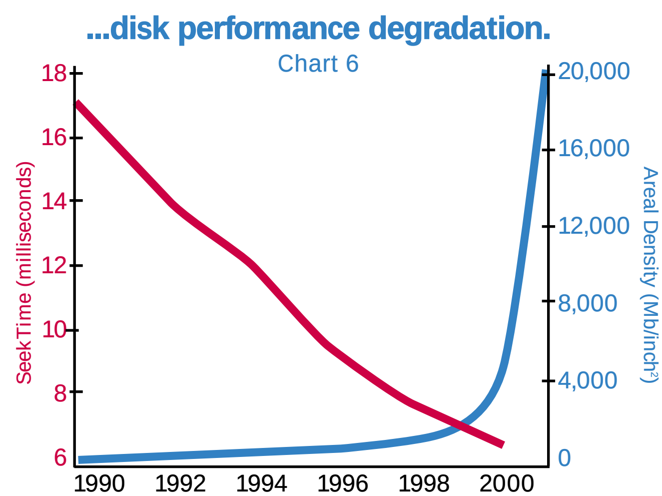

multi-gigabit rates could not be built with off the shelf components. Dr. Barrall and his team would create all server components from scratch. Born was the first Silicon Server architecture, operating at wire speed between gigabit Ethernet at the network side and gigabit fibre channel at the storage end, with drastically fewer components and vastly higher performance.

Suggestive of **EZ Chip**'s (LNOP) pipelined chip architecture, with TOPs (task optimized processors) each designed to perform one specialized function, BlueArc's Silicon Server architecture separates and pipelines the four key server functions into four parallel processing modules. Designed more like a switch or a router than like an ordinary server, the BlueArc device blasts open the bottleneck between users and storage by putting the TCP/IP stack, the file servers, and storage protocols, with all their hundreds of lookup tables, directories, and connection lists, in hardware. One of the parallel modules contains the world's fastest all-hardware TCP/IP networking subsystem. Another holds the world's fastest fibre channel controller. Sandwiched in the middle are diverse protocol translation units and file subsystems that perform all the key storewidth functions–SAN, NAS, you name it–of delivering data from storage to the Web.

BlueArc initially offers a single box with throughput of 2.5 gigabits per second between gigabit Ethernet at one end and gigabit fibre channel at the other. Scaling up to 200 terabytes of storage through a single Ethernet connection, BlueArc blows away the sixteen terabyte limit of current NAS and SAN devices. Consuming nearly half the power and scaling to tens of thousands of simultaneous users, BlueArc pulverizes the current server limits of less than two hundred. For 20 percent of the capital costs, this architecture offers five times the bandwidth for streaming applications, 10 times the bandwidth for traditional file transfers, 30 times the amount of accessible storage, and a hundred times the number of simultaneous user connectivity, all with five-9s solid state reliability. It is a “Silicon Server for the optical age.”

Think of the Silicon Server as a data flow processor between the time-based world of storage and the space-based world of the fibersphere. To the network, it looks like any other Network Attached Storage device. With the traditional server bottlenecks removed it achieves multi-gigabit throughput. It supports all relevant storewidth protocols and can be dropped in to replace any other device that was previously occupying that space. Ultimately gone are all the software based Von Neumann network appliances, whether from **NetApp** (NTAP), **EMC** (EMC), **Procom** (PRCM), **Hewlett Packard** (HWP), or Compaq. With unified loads of 200 terabytes and tens of thousands of clients, gone too are all the load-balancing gear from Cisco, Alteon, **Extreme Networks** (EXTR) and other stalwarts of the network storage center.

BlueArc’s secret is the paradigm. In the past, with protocols, file systems, backup schemes, and I/O technologies in constant change and turmoil, all such functions had to be performed in software. BlueArc recognized that the latest generation of configurable field programmable gate arrays (FPGA) could execute these same tasks in hardware with orders of magnitude improvements in speed, and still be reconfigured in less than one minute, even milliseconds, when new technologies emerged. FPGAs have been a favored *GTR* technology from the beginning and are now being celebrated and explained in depth in our new *Dynamic Silicon* letter written by Nick Tredennick. Represented on our Telecosm list by **Xilinx** (XLNX), **Atmel** (ATML) and now by **Altera** (ALTR) as well, these configurable devices feed on another of the time-space differentials that will shape the future of the industry.

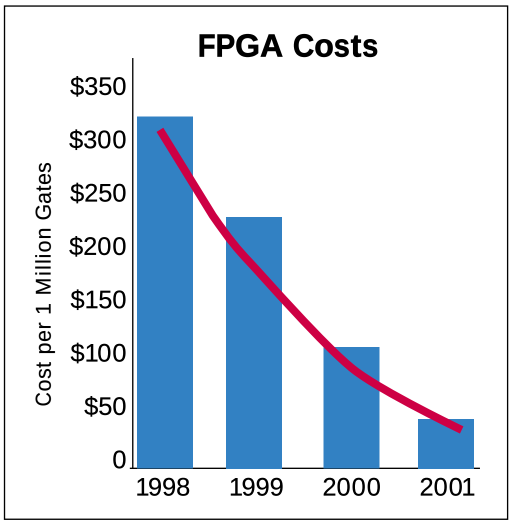

Altera’s storewidth link

Microprocessors tend to economize on space, and waste time. The actual processor occupies less than one tenth of the device; yet the standard von Neumann architecture dictates that the chief way to augment performance is to use this same small silicon area more and more millions of times a second to execute a series of instructions on data fetched for it from off-chip memories. Yet even as this processing time remains precious and scarce, processing “space”–the density of circuits on a chip–has become relatively abundant.

Between 1985 and 2000 circuit density rose one thousand fold while processor speed increased by a factor of twenty. Field programmable gate arrays exploit this new abundance by distributing processing functions across millions of gates operating in parallel across a chip. Such parallelism is usually feasible only when the processing functions are hard wired, limiting flexibility, but FPGAs can be “rewired” in milliseconds. With Xilinx and Altera introducing arrays with as many as 10 million gates or logic elements (comprising some 100 million transistors) this year, increased chip density makes the FPGA an effective rival to the microprocessor for many applications that do not entail a strict temporal order of execution.

Since a storage server primarily directs the data flow, a spatial function, rather than processes it, a temporal function, FPGAs are the natural solution. Just as optical switches outperform electronic switches by merely steering data paths through space rather than processing them in time, parallel FPGAs use millions of gates to steer data across the chip rather than forcing it into time queues to be processed in the CPU.

In the BlueArc box, which uses Altera devices, FPGAs perform functions in the spatial domain, such as spatial data flow, that were previously forced to conform to a temporal processor. Altera Apex FPGAs perform all hardware data transport with mere hundreds of CPU MIPS (millions of instructions per second) rather than the thousands consumed by software inside of traditional servers. FPGAs give the Silicon Server a true performance advantage, one that would not have been possible a year ago and will continue to grow over time, as costs continue to decline.

FPGAs excel in functions–encryption/decryption, pattern recognition, and real time signal processing–that are critical on the Net. In order to move storage to the network, it must be encrypted. In order to communicate across the network, signals must be constantly processed in real time. Speech recognition, optical character recognition, error correction, code division communications, and image capture all entail pattern matching. Transforming the file server into a bi-directional data flow engine, BlueArc is the first company fully to exploit the spatial resources of field programmable gate arrays and channel all data flow inside the hardware.

In the presence of *Dynamic Silicon* author Tredennick, BlueArc ran its first generation Silicon Server through such diverse performance metrics as Iometer (Intel storage benchmark), Spec SFS (UNIX), and Netbench (Windows). “They let me look into the box. They let me change the parameters any way I wanted. Nothing stopped the machine from working at its peak rates.” Not a general purpose server, it does not execute the thousands of server functions, from searches to messaging to graphics rendering that any general purpose Unix or Solaris box from **Sun** (SUNW) or NT machine from Compaq can perform. But BlueArc is the first networked storage appliance that addresses the show-stopping throughput bottleneck imposed by traditional server architectures. It breaks through the storage bottleneck on the Web and reconciles the time domain of the Web with the outer space of the fibersphere. It is a machine for the storewidth warp. We put them on the list both as a potential private investment and a likely public issue.

Mango magic

Virtual private networks (VPNs) are dedicated and encrypted paths that enable security conscious corporations to send mission critical files over shared data networks such as the Internet. **Mangosoft** (MNGX.OB) does not belong to the VPN Consortium, nor does Mango manufacture VPN hardware. Mango does, however, have Mangomind, a LAN-style file service that enables VPN-like applications across the Internet. Mangomind extends the client's existing network services by allowing secure storage and sharing of data between members of extended business workgroups. While traditional VPNs require dedicated phone lines and dedicated IT staff to collect user profiles of remote peering partners and to install, configure, and manage VPN hardware, Mango was able to set Gilder Publishing up with our very own VPN and remote Internet storage cache with a single e-mail.

The first member of a Mangomind drive can invite others to join by e-mailing client software for downloading on their PC. Each member contributing files and folders to the drive can regulate access of other drive members to them. Multiple users can simultaneously and securely access and edit Mind drive files. Use requires no training, no additional hardware, and no IT intervention. The Mind drive does all of the thinking, arbitrating file reads and writes. All applications operate directly on files without creating multiple incoherent copies.

Files saved to the Mango drive are encrypted at the end user prior to transmission over the Internet to Mango's centralized network cache. Managed by **Storage Networks** (STOR) and located at the Exodus data center in Boston, the Mango cache stores data in encrypted form. By contrast, a traditional VPN scheme decrypts files at the destination storage server. In the Mango system, all decryption takes place at the client–on the edge.

Hollowing out the computer by combining the storage and memory of multiple users, Mangomind uses a distributed shared memory architecture originated for massively parallel supercomputers at the late lamented Kendall Square Research. Under the guidance of KSR co-founder Steve Frank, that company developed systems to share memory among many microprocessors within its proprietary supercomputer. Frank then brought the basic concepts to Mango, where they could be executed in software on the accelerated hardware of 2001. Under the guidance of CTO Scott Davis, the company has launched two products–Cachelink, which uses the browser cache in your computer to accelerate access to favored web pages in a company, and now Mangomind, which provides storewidth file services across the Internet.

With the ultimate edge advantage, close proximity to the end user, Mango has accomplished what highly distributed content delivery networks have attempted for years. Mango is as close to the edge of the network that you can get–on the users' desktops. Mango has effectively transformed the desktop into the smartest, most efficient cache on the network, a cache that knows exactly what to cache, when, and for how long. Cache coherency becomes a non-issue. Files are cached and content verified using Mango's version check protocol each time an end user opens or saves a file. Files that have been altered by users working offline are transparently synchronized when the user reconnects to the Net. No extra time, IT power, or CPU power is wasted maintaining cache content freshness.

Hollowing out the computer, Mangomind uses a distributed shared memory architecture originated for massively parallel supercomputers

Most storewidth companies pay top dollar for storage IT professionals to maintain network efficiency and cache coherency on their content delivery and storage networks. Mango, in an ironic twist, doesn't pay, but is paid by its systems administrators. Given that close to 90 percent of the business world is proficient in Windows and that Mangomind was designed with bug-for-bug Windows compatibility, the end user and the cache administrator are one in the same. Although Mango is a very thinly traded stock, we put it on our list with the usual cautions. Among the board members is Nick Tredennick of *Dynamic Silicon*.

– George Gilder with Mary Collins
February 13, 2001

Inserted with this issue is an illustration of *The Storewidth Landscape*. On the back of the insert you will find a list of players in the storewidth space. Companies highlighted are participating in Gilder Publishing's Storewidth 2001 conference.

This month the *Gilder Technology Report*'s Telecosm Table has made the change from fractions to decimals. Also note that technology category titles have changed, but companies within those categories have not changed.

TELECOSM TECHNOLOGIES

ASCENDANT TECHNOLOGY	COMPANY (SYMBOL)	REFERENCE DATE	PRICE	JAN '00: MONTH END	52 WEEK RANGE	MARKET CAP
FIBER OPTICS						
Wireless, Fiber Optic Telecom Chips, Equipment, Systems	Lucent (LU)	11/7/96	11.78	18.60	12.19 - 70.66	63.0B
Wireless, Fiber Optic, Cable Equipment, Systems	Nortel (NT)	11/3/97	11.50	38.23	29.00 - 89.00	116.8B
Optical Fiber, Photonic Components	Corning (GLW)	5/1/98	13.64	56.71	42.83 - 113.33	51.8B
Wave Division Multiplexing (WDM) Components	JDS Uniphase (JDSU)	6/27/97	3.63	54.81	37 - 153.42	52.7B
Adaptive Photonic Processors	Avanex (AVNX)	3/31/00	151.75	64.63	39 - 273.50	4.2B
All-Optical Cross-Connects, Test Equipment	Agilent (A)	4/28/00	88.63	54.55	38.06 - 162.00	24.9B
Tunable Sources and WDM Components	New Focus (NUFO)	11/30/00	20.31	60.19	16 - 165.13	3.8B
Crystal-Based WDM and Optical Switching	Chorum (private)	n/a	n/a	n/a	n/a	n/a
WDM Metro Systems	ONI (ONIS)	12/29/00	39.56	55.50	22.25 - 142	7.3B
LAST MILE						
Cable Modem Chipsets, Broadband ICs	Broadcom (BRCM)	4/17/98	6*	109.94	72.38 - 274.75	25.9B
S-CDMA Cable Modems	Terayon (TERN)	12/3/98	15.81	6.50	3.50 - 142.63	428.3M
Linear Power Amplifiers, Broadband Modems	Conexant (CNXT)	3/31/99	13.84	18.06	12.63 - 132.50	4.4B
WIRELESS						
Satellite Technology	Loral (LOR)	7/30/99	18.88	5.90	2.69 - 21.00	1.8B
Low Earth Orbit Satellite (LEOS) Wireless Transmission	Globalstar (GSTRF)	8/29/96	11.88	0.81	0.72 - 36.44	86.1M
Code Division Multiple Access (CDMA) Chips, Phones	Qualcomm (QCOM)	7/19/96	4.75	84.06	51.50 - 162.56	63.3B
Nationwide CDMA Wireless Network	Sprint (PCS)	12/3/98	7.19 *	30.50	17.63 - 66.94	28.4B
CDMA Handsets and Broadband Innovation	Motorola (MOT)	2/29/00	56.83	22.81	15.81 - 61.54	49.8B
Wireless System Construction and Management	Wireless Facilities (WFII)	7/31/00	63.63	40.88	27.13 - 163.50	1.8B
GLOBAL NETWORK						
Metropolitan Fiber Optic Networks	Metromedia (MFNX)	9/30/99	12.25	15.94	9.13 - 51.88	8.8B
Global Submarine Fiber Optic Network	Global Crossing (GX)	10/30/98	14.81	22.02	11.25 - 61.81	19.5B
Regional Broadband Fiber Optic Network	NEON (NOPT)	6/30/99	15.06	16.31	3.50 - 159.0	305.9M
Telecommunications Networks, Internet Backbone	WorldCom (WCOM)	8/29/97	19.95	21.56	13.50 - 52.50	62.2B
Global Submarine Fiber Optic Network	360networks (TSIX)	10/31/00	18.13	14.50	10.0 - 24.19	11.8B
STOREWIDTH						
Directory, Network Storage	Novell (NOVL)	11/30/99	19.50	8.63	4.78 - 44.56	2.7B
Java Programming Language, Internet Servers	Sun Microsystems (SUNW)	8/13/96	6.88	30.56	25 - 64.66	98.4B
Network Storage and Caching Solutions	Mirror Image (XLA)	1/31/00	29	9.03	2.81 - 112.5	957.9M
Disruptive Storewidth Appliances	Procom (PRCM)	5/31/00	25	21.31	10.25 - 89.75	247.6M
Remote Storewidth Services	Storage Networks (STOR)	5/31/00	27*	27.94	16.50 - 154.25	2.6B
Complex Hosting and Storewidth Solutions	Exodus (EXDS)	9/29/00	49.38	26.63	14.88 - 89.81	14.7B
Hardware-centric Networked Storage	BlueArc (private)	1/31/01	n/a	n/a	n/a	n/a
Virtual Private Networks, Encrypted Internet File Sharing	Mangosoft (MNGX.OB)	1/31/01	1.00	1.00	0.75 - 28.0	26.9M
MICROCOSM						
Analog, Digital, and Mixed Signal Processors	Analog Devices (ADI)	7/31/97	11.19	62.60	42.63 - 103.00	22.4B
Silicon Germanium (SiGe) Based Photonic Devices	Applied Micro Circuits (AMCC)	7/31/98	5.67	73.59	32.74 - 109.75	21.8B
Programming Logic, SiGe, Single-Chip Systems	Atmel (ATML)	4/3/98	4.42	17.00	9.38 - 30.69	7.9B
Single-Chip ASIC Systems, CDMA Chip Sets	LSI Logic (LSI)	7/31/97	15.75	24.75	16.30 - 90.39	7.9B
Single-Chip Systems, Silicon Germanium (SiGe) Chips	National Semiconductor (NSM)	7/31/97	31.50	28.70	17.13 - 85.94	5.0B
Analog, Digital, and Mixed Signal Processors, Micromirrors	Texas Instruments (TXN)	11/7/96	5.94	43.80	35.00 - 99.78	75.8B
Field Programmable Gate Arrays (FPGAs)	Xilinx (XLNX)	10/25/96	8.22	54	35.25 - 98.31	17.8B
Seven Layer Network Processors	EZchip (LNOP)	8/31/00	16.75	16.44	5.63 - 43.75	106.1M
Network Chips and Lightwave MEMS	Cypress Semiconductor (CY)	9/29/00	41.56	27.33	18.25 - 58.00	3.6B
Field Programmable Gate Arrays (FPGAs)	Altera (ALTR)	1/31/01	30.25	30.25	19.63 - 67.13	12.0B

ADDED TO THE LIST: BLUEARC, MANGOSOFT, AND ALTERA **DELETED FROM THE LIST: C-CUBE** * INITIAL PUBLIC OFFERING

NOTE: The Telecosm Table is not a model portfolio. It is a list of technologies in the Gilder Paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. Mr. Gilder and other GTR staff may hold positions in some or all of the stocks listed.

Gilder Technology Report ***Published by Gilder Publishing, LLC and Forbes Inc.***

Monument Mills • PO Box 660 • Housatonic, MA 01236
Tel: (888)484-2727 • Fax: (413)274-3030 • Email: info@gildertech.com

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

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

FOR SUBSCRIPTION INFORMATION TELEPHONE TOLL FREE:

(888) 484-2727

Website:
www.gildertech.com