GILDER TECHNOLOGY REPORT

March 2003 / Vol. IX No. 3

Memory, Storage, & Untethered Devices

Familiar PC components —flash memory, SRAM, DRAM, and the hard disk— won't do for emerging untethered systems.

Inside:

- Semiconductor memory
- MEMS-based storage example: IBM Millipede
- PMCm—A New Non-volatile Memory
- The hard disk
- The personal computer
- The future
- The transistor and the microprocessor
- The right stuff
- Lessons

Fifteen hundred contacts grace my Microsoft Outlook address book. My PDA synchronizes with Outlook, so I have phone numbers and e-mail addresses wherever I am. Neither the PDA nor Outlook synchronizes my cell phone, so I enter those numbers separately. Some cell phones do this, but they're still too expensive for me. Besides, it wouldn't help much anyway. My address book is always out of date. Some 20% of the entries are wrong—and I don't know which ones. The root of the problem is that I store these contacts locally. Perhaps a thousand people have my address filed. The post office changed my zip code seven or eight years ago, and I still get mail with the old zip code.

The solution: one place for my contact information and a list of people authorized to access it. Instead of having my address in a thousand places, it's in one. If I move, there's one address to change. That beats notifying everyone who I think has my address, and it beats expecting them to update their local copies.

The story highlights the tradeoff between bandwidth and storage. If bandwidth were plentiful, my PDA and cell phone wouldn't need local storage for my address book or for anything else. They would rely on their connection to source current information. But high end-to-end bandwidth is scarce. Scarce bandwidth extends storage. Whatever life span storage would have if high bandwidth were the norm is extended by bandwidth's absence. Similarly, low bandwidth inflates the demand for greater storage capacity. Think of the supply of bandwidth and storewidth as zero-sum elements.

Meanwhile, the PC is riding off into the sunset. Untethered devices are replacing the PC as the next dominant electronic platform. This change of platform venue has new rules declaring our PC-bred electronic components (flash memory, SRAM, DRAM, hard disk) inadequate. New non-volatile memory and storage are to be the key enablers of the new platform.

Terminology: I find it useful to distinguish a place to put information being worked on ("memory") from a general repository ("storage").

Semiconductor memory

Semiconductor memory grew up with the personal-computer business. The consumer-oriented personal computer competed on microprocessor performance. When the personal computer was introduced, the microprocessor and its dynamic random-access memory (DRAM) were about the same speed. Microprocessor makers have always optimized speed, and the memory makers have always optimized capacity. So, over time, the speed of microprocessors and DRAMs diverged; microprocessors became much faster than DRAMs. Since it wouldn't do to have a fast microprocessor waiting on relatively slow DRAM, system makers put small, fast semiconductor memories between the microprocessor and the DRAM to cache frequently used information. This static random-access memory (SRAM) is built with the same fast transistors used in microprocessors, so SRAM can be as fast as the microprocessor. But SRAM is expensive, with only a sixteenth of the capacity of DRAM.

Both SRAM and DRAM lose their information when the power is off, so PCs need a third type of semiconductor memory to hold the startup information needed

CONTINUED ON PAGE 6

MEMS-based storage example: IBM Millipede

The hard-disk platter in Hitachi Global Storage Technologies' Microdrive is the size of a U.S. quarter. But the Microdrive is still too power hungry for untethered devices. How about shrinking it even more? Perhaps it can be built on a chip. That's what microelectromechanical systems (MEMS)—chips with moving parts—are all about. The industry needs bulk storage for untethered devices, but a miniature hard disk isn't the answer. The rotating platter burns too much power and starting and stopping it, to save power, takes too much time.

HITACHI GLOBAL STORAGE TECHNOLOGIES

Fig. 1. The 4-GB version of Hitachi Global Storage Technologies' Microdrive is due in late 2003.

Bad news. Developing bulk storage for untethered devices means leaving fifty years of experience behind in building hard disks. It means inventing new ways to store and to retrieve data. Inventing techniques takes more time than making incremental improvements. MEMS-based storage won't burst onto the scene today and take over the world tomorrow. It will take years. Inventing bulk storage for untethered devices means experimentation, false starts, competing solutions, and developing standards. The changing requirements and new production methods open a host of options. The range of MEMS-based storage approaches will recall the early days of integrated circuits. Anyone remember integrated circuit families (TTL, ECL, DCTL, etc.) and circuit types (NMOS, PMOS, CMOS, bipolar, etc.)?

IBM's Zurich research division has experimented with MEMS-based storage for years. In late 2002, IBM announced improvements. While the company's "Millipede" storage is years from commercialization, it illustrates the decisions to be made in developing bulk storage for untethered devices.

Two postage-stamp-sized chips comprise Millipede.

The stationary chip is an array of read/write probes. From above, each probe looks like a rounded "v," is attached by its arms, and has a point at the end like a phonograph needle. The chip also contains read and write circuits for the probe tips, position sensors, permanent magnets, and coils. The sensors, magnets, and coils are part of the electromagnetic actuator circuits that move the second chip in the x, y, z, and tilt directions. Fig. 3 shows part of the read/write probe array and the details of a single probe and its tip.

Fig. 2. A section of the IBM Millipede MEMS-based storage chips.

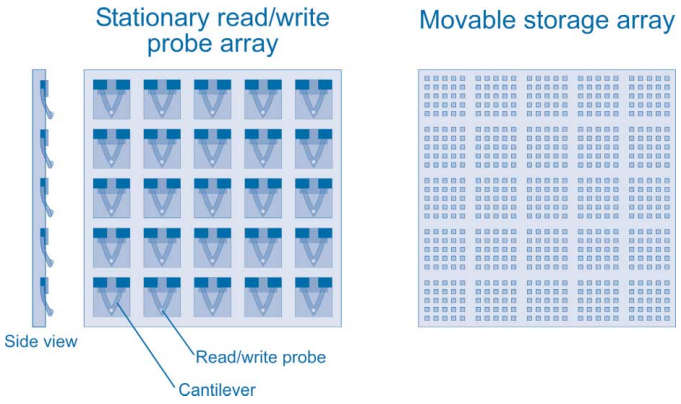

Fig. 3. (Left) Top view of read/write probes arrayed on the IBM Millipede MEMS-based storage chip. (Right) Side view close up of a single probe at rest on the storage medium.

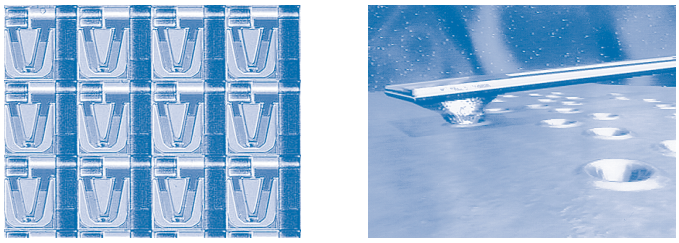

Fig. 4. 1,024 read/write probes occupy the center 3mm x 3mm (dark area) on the Millipede chip.

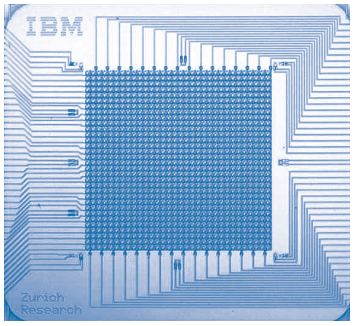

The movable chip is the storage medium. The medium is two layers of polymer over a silicon base. Bits are stored in the soft top-layer polymer. The middle-layer polymer, which is harder than the soft polymer but is softer than the silicon substrate, protects the probe needles from wear against the silicon.

Millipede uses thermomechanical storage. Storing a bit means running electrical current through a probe, heating its tip to 400 °C. The hot tip melts a tiny crater in the top-layer polymer. Reading a bit heats the probe to 350 °C (not hot enough to melt the polymer). If the probe is resting in a mini-crater, it runs cooler and more current flows. If the probe is not in a mini-crater, it runs hotter and less current flows. The current flow betrays the presence or absence of mini-crater "bits." A different melting process erases individual bits.

Electromagnetic actuators move the storage-medium chip, over the 32x32 array of probes, in x and y directions to build the rows and columns of data. The storage area looks

like a 32x32 checkerboard where a square in the checkerboard contains a million bits. Each probe only reads and writes the bits in its square.

IBM's experimental chip stores 0.9 GB and reads and writes data at 32 kb/s. That's not impressive for density or for performance. But, based on this experiment, IBM believes that it can achieve densities of 1,000 Gb/in^2 with much faster data rates. (Data rates in this experiment were limited by the PC exchanging data with the chip.) This 1,000 Gb/in^2 compares well with today's leading storage densities of 50 Gb/in^2.

Millipede has 1,024 read/write probes, compared to one per platter for hard disks. In hard disks, read/write probes are expensive, so the fewer the better. MEMS-based probes are batch fabricated in silicon; the cost to build them by the thousands is insignificant. But a Millipede structure would mean big changes in how data is organized on the medium.

Millipede uses electromagnetic actuators to move the storage medium. Other MEMS devices use electrostatic actuators or use linear drive motors. Electromagnetic actuators require bulky coils and bulky permanent magnets, and they require continuous current to maintain position. IBM's storage medium is a soft polymer. IBM has experimented with polymers and thinks there's room for improvement. Other MEMS-storage devices use magnetic media, phase-change media, and even the positioning of individual atoms. Millipede's probes contact the storage medium; other systems use non-contact probes. Contact probes wear out faster, but non-contact probes may require significantly more energy to operate.

Every aspect is unresolved. Slogging through alternatives slows adoption. IBM has done the most advanced MEMS-based storage experiments that I'm aware of, but if I had to choose winners, I'd bet on electrostatic actuators and on magnetic storage for the short term (three years) and on some version of atomic storage for the long term (ten years).

PMCm—A New Non-volatile Memory

Programmable Metalization Cell memory (PMCm) is a fresh approach to non-volatile memory from the field of solid-state electrochemistry. Startup **Axon Technologies Corporation** (www.axontc.com), founded in 1996, is the creator and licensor of PMCm. PMCm uses the principle of an electrolyte. Electrolytes use ions (electrically charged atoms) as the means for electric current flow. It's how most batteries work. Atoms at the positive terminal give up an electron and become positively charged ions. The ions move through the electrolyte toward the negative terminal. In car batteries, the electrolyte is sulfuric acid; in "dry cell" batteries, the electrolyte is a paste. PMCm uses thin (a few nanometers), solid electrolyte sandwiched between a silver terminal and a conductor (e.g., copper) terminal. Here's how this ingenious structure works.

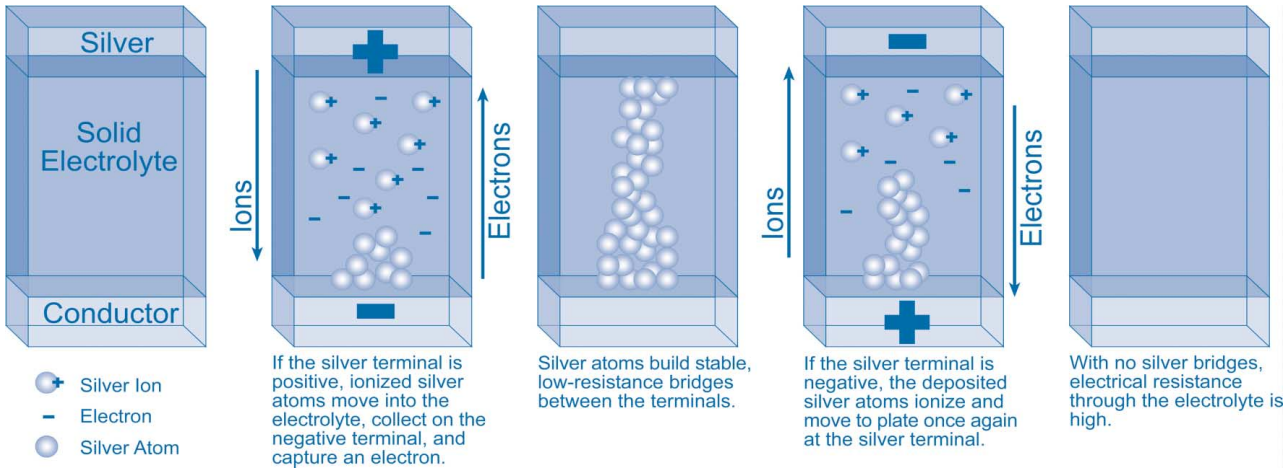

Fig. 5. Programmable Metalization Cell memory builds conducting paths between the device's terminals in a reversible process that changes electrical resistance by orders of magnitude.

A small positive voltage, ~0.3 volts, on the silver terminal creates silver ions that migrate into the electrolyte toward the negative terminal. Bits are stored by the presence or absence of the silver conducting paths. Measuring electrical resistance between the terminals reveals whether a one or a zero is stored. Changing the electrical polarity reverses the process, ionizing the silver atoms in the electrolyte and causing them to migrate back to the silver terminal. Switching times can be under 10 nanoseconds.

Because there are no floating-gate charges (as in semiconductors), no magnetic or electric dipoles to coerce, and no heating elements, this process operates at low voltage, with low energy requirements, and at ambient temperatures. Since there are no appreciable electric, magnetic, or thermal fields associated with the bit cell, the cells can be closely packed without worry of crosstalk. The process is reversible, written values are non-volatile, and the cells do not wear out through trillions of cycles. In addition, cells can be as small as half the size of a DRAM cell; cells can be added with metallization layers (after the rest of the circuit is built); and the cells are built using a standard semiconductor manufacturing process.

TELECOSM TECHNOLOGIES

Ciena	(CIEN)
Corvis	(CORV)
JDS Uniphase	(JDSU)
Avanex	(AVNX)
Essex	(ESEX.OB)
Equinix	(EQIX)
Sprint PCS	(PCS)
Qualcomm	(QCOM)
Broadcom	(BRCM)
Altera	(ALTR)
EZchip	(LNOP)
Terayon	(TERN)
National Semiconductor	(NSM)
Synaptics	(SYNA)
Intel	(INTC)
Soma Networks	(PRIVATE)
Narad Networks	(PRIVATE)
Flextronics	(FLEX)
Taiwan Semiconductor	(TSM)
Transmeta	(TMTA)
Analog Devices	(ADI)
ARC Cores	(ARK.L)
ARM Limited	(ARMHY)
Calient	(PRIVATE)
Celoxica	(PRIVATE)
Cepheid	(CPHD)
Chartered Semiconductor	(CHRT)
Coventor	(PRIVATE)
Cypress	(CY)
Cyrano Sciences	(PRIVATE)
Energy Conversion Devices	(ENER)
Foveon	(PRIVATE)
Legend Group Limited	(LGHLY.PK)
Microvision	(MVIS)
QuickSilver Technology	(PRIVATE)
SiRF	(PRIVATE)
Tensilica	(PRIVATE)
Triscend	(Private)
United Microelectronics	(UMC)
VIA Technologies	(2388.TW)
Wind River Systems	(WIND)
Xilinx	(XLNX)

Ciena (CIEN)

METRO WDM PLATFORMS

MARCH 17: 4.88 52-WEEK RANGE 2.41-9.62 MARKET CAP: 2.1B

Sales of $70.5 million for 1QFY03 came in above expectations. Ciena reported nine new customers and recognized revenue from a total of 64 customers, up from 58 in Q4. Strength in the quarter came from the two divisions that have weighed most heavily on the company during this downturn: long haul DWDM and metro. Stabilization, hinting at the beginning of a recovery, was seen in long haul as sales increased over 61% q/q. Going forward, we look to the continued rumors surrounding long-haul opportunities with AT&T, the government, and several European carriers. Metro reversed five consecutive quarters of market-share losses with q/q sales growth of over 100%, reaching 22% of total sales for the quarter. Evidence that Ciena is beginning to benefit from its acquisition of ONI is a topic that we hope to address further when Rohit Sharma, ex-CTO of ONI, speaks at this month's Gilder/Forbes Storewidth conference.

Corvis (CORV)

WDM SYSTEMS, RAMAN AMPLIFICATION, EDGE SWITCHES

MARCH 17: 0.57 52-WEEK RANGE 0.47-1.46 MARKET CAP: 229M

Broadwing announced the sale of its broadband business to C III communications, a new joint venture between Corvis and Cequel III, a venture-capital firm backed by Jerry Kent, the founder of cable firm Charter Communications, Inc. Broadwing will receive $129 million cash, and C III will assume all current liabilities at the IXC level, which includes $375 million of long-term debt. More importantly, C III will not assume the $46 million, 9% notes; the $451 million, 12.5% preferred stock; or IXC's $155-million credit facility.

Corvis is responding to Clayton Christensen's insights on integration and modularity. At a time of rapid technological change, integration is imperative. There is no time to standardize all the interfaces of the system. Corvis is pursuing a bold strategy of expansion rather than waiting around for the industry to turn up again. Its technology will enable the company to win any price war for long-distance data transmission.

Essex (ESEX.OB)

OPTICAL PROCESSORS

MARCH 17: 3.10 52-WEEK RANGE 1.50-6.40 MARKET CAP: 24M

Essex acquired privately owned Sensys Development Laboratories, Inc. (SDL), a telecomunications company that provides systems and software engineering services, and is expected to add $4 million to Essex's 2003 revenues.

Sprint PCS (PCS)

NATIONWIDE CDMA WIRELESS NETWORK

MARCH 17: 4.00 52-WEEK RANGE 1.75-13.45 MARKET CAP: 3.9B

One of several major issues facing wireless operators is wireless number portability (WNP), the ability to switch wireless carriers while retaining your original wireless number. The Big Six have uniformly stated their general opposition to WNP, weary of a significant increase in churn. Merrill Lynch identifies the following as key risk factors in distinguishing the potential winners and losers of WNP: the size of the carrier's subscriber base; the number of business customers a carrier has; and customer perception of coverage and quality, including the implications of technology migration. Focusing, as you might expect, on the issues of coverage quality and technology migration, we view WNP as an opportunity for the CDMA-based operators, Sprint PCS and Verizon Wireless, to emerge as clear winners. CDMA2000 allows these carriers the ability to deliver content at 2-3 times the speed offered by the GSM-based operators. More importantly, as voice remains the wireless killer app, we return to the fact that GSM-based data networks, be it GPRS or EDGE, degrade voice capacity, while CDMA2000 enables a near doubling. Sprint PCS continues to deliver the clearest Vision of how wireless voice and data connectivity can dramatically increase productivity. However, we fear that the perception of poor network coverage as well as sub-par customer support systems threaten to undercut the quality of its technology offerings. The company has focused on these areas in their last two conference calls and must continue to do so.

Qualcomm (QCOM)

CDMA INTEGRATED CIRCUITS, IP, SOFTWARE

MARCH 17: 37.57 52-WEEK RANGE 23.21-44.00 MARKET CAP: 30B

Qualcomm revealed that its new best friend China Unicom will begin the world's first GSM1x commercial trial. The trial, which is slated for the second half of this year, will proceed in several stages, the first of which will be a demonstration of CDMA2000 data service with a GSM core network. This deployment will also provide a real-world environment to perform testing of dual-mode GSM/CDMA2000 handsets. China Unicom is currently running two networks: its original GSM-based network with 60-million subscribers, as well as its highly publicized, one-year-old CDMAOne, soon to be (end of March) a CDMA2000-based network with over 8-million subscribers. GSM1x technology will give Unicom the ability to offer its 60-million GSM users all

the benefits of the CDMA2000 air interface—spectral efficiency, increased voice capacity, and increased data rates—without abandoning the company's investment in its reliable and widely deployed GSM-MAP (mobile application part). The potential benefit to Qualcomm is clear—60-million chipsets clear.

Broadcom (BRCM)

BROADBAND INTEGRATED CIRCUITS

MARCH 17: 16.00 52-WEEK RANGE 9.52-40.62 MARKET CAP: 4.4B

Broadcom shipped over 1-million 802.11g chipsets in the last quarter, impressing the Street. However, the much more pristine 5-GHz spectrum used by 802.11a piqued our interest with the news that Atheros Communications won placement with IBM, HP, NEC, and Toshiba to incorporate the company's 2.4-GHz and 5-GHz spectrum chipsets in its laptops, providing 802.11a/b multimode connectivity.

Altera (ALTR)

PROGRAMMABLE LOGIC DEVICES

MARCH 17: 13.84 52-WEEK RANGE 8.32-24.66 MARKET CAP: 5.3B

Xilinx reiterated its revenue estimates for the upcoming March quarter, further indicating expectations toward the high-end. Strength in the quarter was outlined in general terms: Asia-Pacific and Europe, strong; North America, flat. Altera was more forthcoming with its announcement of lifted expectations. The company reported continued strength in its Stratix family products, as well as a strong ramp for its Cyclone FPGA. The Cyclone, which began shipping in December, now claims more than 225 customers, up from 75 in January.

National Semiconductor (NSM)

ANALOG EXPERTISE, FOVEON IMAGERS

MARCH 17: 17.85 52-WEEK RANGE 9.95-37.30 MARKET CAP: 3.2B

February 20 saw National Semiconductor unveil elements of a "road map" for enhanced long-term profitability. Details included cost reductions, divestitures, and outsourcing. The company has put up for sale two of its least profitable businesses: digital baseband products for GSM, GPRS, and WCDMA cellular handsets; and its information appliance unit based on the Geode family of processors, the last remnants of the fruitless Cyrix Semiconductor acquisition. National no longer feels that it is necessary to have a digital baseband solution in order to gain content in mobile handsets. Management points to its partnership with ARM, whereby National develops power-management products and standard interfaces with ARM cores. This view is further confirmed by the fact that seven of the company's top-ten customers are handset manufacturers. National increases its content within handsets through analog products for power management, audio, color displays, and cameras. Handset content for the company today is around $3-$4 plus another $6-$8 for color solutions, compared to around $1 just a few years ago. On the outsourcing side, National signed an agreement with Taiwan Semiconductor for all logic running on a 150-nanometer-and-below process technology. This agreement implies that National has freed itself from the capital expense of building fabs for leading-edge, next-generation processes. This move toward horizontal fragmentation improves National's ability to deliver "value transistors" to the consumer-centric OEMs that make up the majority of the company's customers. Horizontal fragmentation also increases the efficiency of National's engineers.

> National reported FY3Q03 revenues of $404 million, down 4.7% q/q. Positively, bookings in the quarter increased 3%, while analog bookings increased 12%, and portable power management bookings increased 50%. This higher Q4 backlog has lead management to project Q4 revenues of $420-$432 million, up 4% to 7%.

Synaptics (SYNA)

TOUCH-SENSORS, FOVEON IMAGERS

MARCH 17: 7.24 52-WEEK RANGE 3.13-20.75 MARKET CAP: 171M

As Symbian, a developer of advanced mobile phone operating systems battles it out with Microsoft, among others, in what is being called "The Next Browser War," Synaptics has signed on as a member of its Platinum Partner Program. Synaptics will develop a version of its Spiral pen input solution to run on the Symbian OS. The Spiral's inductive position-sensing technology enables the position of the pen to be measured above an LCD while the sensor board lies underneath the display. This position significantly increases the display quality by eliminating the need for a screen overlay. In addition, elimination of the touch screen reduces the device's backlighting requirements, substantially reducing battery consumption.

Intel (INTC)

MICROPROCESSORS, SINGLE-CHIP SYSTEMS

MARCH 17: 17.17 52-WEEK RANGE 12.95-32.25 MARKET CAP: 113B

As the rumblings within the Intel Communications Group grew louder, we took notice, even resigning ourselves to the possibility that the "Intel of the Telecosm" may in fact turn out to be Intel itself (see *GTR*, September 2002). The March 12 launch of its Centrino mobile PC processor turned Intel's rumblings into roars. Intel intends to spend $150 million this year on Wi-Fi-related companies around the world. So far, its focus has largely been on subsidizing the public Wi-Fi access market. We worry about the glorified status these "hot spots" are garnering and continue to question the commercial viability of such networks, believing, rather, that desktop replacement notebooks, which can take advantage of WLAN connectivity at home, the office, airport lounge, or hotel room, will drive uptake. Intel plans to translate its success in the desktop space into notebooks, clearly indicating that any and all OEMs planning to take advantage of Intel's brand name and marketing program must use the Banias processor, the Odem/Montara graphics chipset, and the Calexico WLAN chipset. Estimates claim that these requirements will yield Intel a 20% share of PC WLAN revenue (excluding access-point chipsets) heading into 2003, ramping to 50%+ in 2004 as the company launches its 802.11a/b solution in June, followed by its 802.11a/b/g solution some time in 2H03.

> Flashback: Intel's attempt to raise prices on flash memory by 20%-40% beginning in January of this year appears to have backfired. CIBC believes that several of the company's largest handset customers, some 80%+ sourced to Intel, have been actively designing-in competitive products in an attempt to decrease their dependence on Intel over the course of 2003.

Analog Devices (ADI)

RF ANALOG DEVICES, MEMS, DSPS

MARCH 17: 28.76 52-WEEK RANGE 17.88-46.82 MARKET CAP: 10.4B

Strong sales from the communications division, primarily wireless handsets, base stations, and broadband access (DSL) markets allowed Analog Devices to report better than expected sales of $467 million. Geographically, sales to Asian and European customers came in stronger than North America and Japan.

ARM Limited (ARMHY)

MICROPROCESSOR AND SYSTEMS-ON-CHIP CORES

MARCH 17: 2.60 52-WEEK RANGE 1.87-13.24 MARKET CAP: 836M

ARM Holdings and Synopsys announced a reference design methodology for integrating synthesizable ARM processor cores into the designers' physical implementation process. Synopsys enables ARM customers to perform their own physical implementation, while at the same time ensuring compliance within ARM's architecture.

CONTINUED FROM PAGE 1

by chips on the system board. Originally, that was read-only memory (ROM), but ROM has been displaced by flash memory. The problem with ROM is that it could never be changed. Flash memory holds information even when power is off, and the information can be changed.

Why not use flash memory for everything? Flash memory is slow compared to SRAM, and it actually wears out over time.

Here's the situation. Flash memory is for information that needs to survive power on-off cycles. It holds the boot information for your personal computer, it holds the pictures in your digital camera, and it holds the telephone numbers in your cell phone. DRAM is the working memory for your personal computer. It holds running programs. Levels of SRAM sit between the microprocessor and the DRAM in an attempt to speed-match the fast processor to the slow DRAM.

The hard disk

The hard disk has dominated storage since 1956. There have been numerous reports of its impending demise. But the hard disk has been amazingly durable.

Pundits said "bit density will soon reach the theoretical limit; bits cannot get smaller and retain data reliably." But soon bits *were* smaller, and the theory was adjusted. Theory and practice have leapfrogged several times.

Enthusiasts said the per-bit cost of semiconductor storage would beat the cost of magnetic storage. "There's no way these electromechanical clunkers can compete with semiconductor storage." But semiconductor storage isn't even close; the hard disk is the hands-down winner. In 2003, the "sweet spot" for hard disks is 120 GB priced at $108. Corresponding flash-memory storage is 256 MB (¼ GB) priced at $15. Even the sweet-spot hard disk (not the biggest) has 500 times the capacity of that flash memory; the per-bit cost of the flash memory is 65 times that of magnetic storage.

The hard disk isn't going away. It'll remain the bulk storage champion in tethered systems. MEMS are invading the hard disk, improving the hard disk's performance and capacity. The PC's system board might even invade the hard disk. Imagine a hard disk with a PC system board, power cord, Ethernet, and a half-dozen USB ports in a package the size of today's hard disk. This portable desktop computer packs easily between home and office. Plug in the monitor, keyboard, and mouse that stay with desks, and you retain your personal environment and all of your data.

The personal computer

The PC made memory and storage what they are today. Flash memory's competence is its non-volatility—keeping data across power cycles. Flash memory holds initialization programs. SRAM's competence is speed. DRAM's competence is large working memory. DRAM is the microprocessor's working memory for the operating system and application programs. The hard disk's competence is bulk storage.

Microprocessor speeds are in picoseconds (one-trillionth of a second); semiconductor memory speeds are in nanoseconds (one-billionth of a second); and magnetic storage speeds are in milliseconds (one-thousandth of a second). Semiconductor memories hold megabytes; magnetic storage holds gigabytes. The PC built the family of memory and storage—flash memory, SRAM, DRAM, and the hard disk. As PC generations improved, each of these components improved and secured its niche. But the world is changing; the value PC signals the end of the PC dynasty.

The future

The world is splitting into tethered and untethered systems. Tethered systems, such as PCs, run from wall power; untethered systems, such as cell phones, run from batteries. Tethered systems constitute the "fibersphere"—the world's collection of computing, data ports, data networks, and storage. Untethered systems connect wirelessly to the fibersphere and are the collectors and consumers of data.

The engineering talent, once devoted to improving the PC, will be redirected to developing untethered systems. The design objective for the PC has been cost performance; the design objective for untethered systems is cost-performance-per-watt. Untethered systems want low cost, excellent performance, and long battery life. The requirements of untethered systems will change the investment emphasis in memory and in storage. The PC's electronic legacy is a set of components well suited to the system structure of the PC. These components are unsuited to emerging untethered systems. They fall short in absolute performance, in power consumption, or both. *The component family—flash memory, SRAM, DRAM, and the hard disk—that has served the tethered PC for over twenty years, won't do for untethered systems.*

The hard disk, for example, is ill suited to untethered devices. While it retains data when power is off, its startup delay is long. It takes seconds for the disk to spin up. Shutting off the hard disk between accesses saves power, but suffers startup delays. Leaving the hard disk spinning burns power continuously, which shortens battery life. Further, a hard disk can be made only so small, and there's a premium for making it small. **Hitachi**'s (HIT) one-inch, 1-GB Microdrive costs twice as much as a current 120-GB hard disk. The largest hard disk in a handheld device is the Toshiba 20-GB hard disk in **Apple**'s (AAPL) iPod. Hard disks for portable devices cost more, they forfeit capacity, and they don't satisfy needs for long battery life or for instant access to information.

SRAMs, like microprocessors, don't match the needs of untethered devices either. SRAMs and microprocessors are built of logic transistors, which burn lots of power in exchange for speed. Portable devices need the speed, but they cannot afford the power.

The transistor and the microprocessor

The transistor is the atom of electronic systems. The differences between a personal digital assistant, a GPS receiver, an MP3 player, and a cell phone are in how their millions of transistors are wired together. Given a budget of, say, a billion transistors, an engineering team could deliver any one of these systems. In the old days, with individually packaged transistors, it would have been possible to disassemble a radio and to use the same transistors to build a voice recorder or to build an autopilot. Now, the transistors are integrated onto chips. Instead of physically unsoldering and soldering transistors, suppose we could change the connections to each transistor after manufacture. We'd load one configuration into the chip to make it a cell phone, another would personalize it as an MP3 player, and so on. Too much overhead. (Today, it would take many bits to configure *each* transistor in the final design.)

What's the alternative if we want the same flexibility, but we cannot afford a large number of configuration bits? We use a microprocessor. Its circuits are fixed at manufacture, but it's flexible because programs characterize its behavior. We don't have the huge number of configuration bits, but inefficiency reigns because any desired behavior can only be expressed with the microprocessor's limited vocabulary (instruction set).

There's plenty of middle ground between configuring individual transistors and making do with a microprocessor's fixed vocabulary. Programmable logic devices (PLDs), whose dominant manufacturers are **Altera** (ALTR) and **Xilinx** (XLNX), are in this middle ground. PLDs intersperse an array of Lego-block circuits in a wire grid. Configuration bits in an on-chip SRAM attach wires to the circuit blocks, combining generic blocks into useful circuits. This way of implementing functions is still much more efficient than the microprocessor's instruction-based approach. And the PLD's configuration overhead is much less than it would be dealing with individual transistors. But the PLD's circuits are still slowed by the SRAM configuration bits and there's substantial overhead in the SRAM itself.

But suppose the Lego-block circuits are specialized for use in cell phones. Suppose these specialized circuit blocks contain thousands or tens of thousands of transistors and that there are fewer than ten such blocks. Now fewer configuration bits can connect more efficient circuits, so configuration overhead drops dramatically and performance improves. That's the idea behind startups such as **Ascenium** and **QuickSilver Technology**.

The right stuff

"Non-volatile" describes memory that holds information when the power is off. Flash memory is non-volatile and hard disks are non-volatile. Untethered devices need non-volatile storage, but neither flash memory nor hard disks are right for them. Flash memory is too slow to read and it is really slow to write. Untethered devices need non-volatile memory to initialize their circuits and to store the operating system and application programs.

Tethered applications use flash memory to initialize their circuits, and they use hard disks for bulk storage. They use DRAM for the operating system and applications, and they use SRAM to bridge the speed gap between the microprocessor and the DRAM. This hierarchy is too elaborate for untethered devices.

Today's untethered devices don't want hard disks. The hierarchy of flash memory, SRAM, and DRAM means too many chips and too much power. SRAM has to go because it's power hungry, but losing the SRAM means that the digital signal processor and the microprocessor's performance-oriented duties have to go as well. Flash memory stays because the portable device needs non-volatility. DRAM stays because today's untethered devices need working memory for the operating system and applications. These programs cannot run from the flash memory because flash memory would wear out in a few hours or a few days. The flash memory in your personal digital assistant and digital camera seem to last forever because you don't read and write these memories tens of millions of times a second, which is what a microprocessor running its programs from flash memory would do.

The emergence of untethered systems creates enormous incentive to displace flash memory, SRAM, DRAM, and the hard disk. The ideal memory has the non-volatility of flash memory, the density of DRAM, and the speed of SRAM. Commercial alternatives to flash memory and to DRAM have been around since **Ramtron** (RMTR) introduced ferroelectric random-access memory (FRAM) in 1988. The prominent candidates are FRAM, magneto-resistive random-access memory (MRAM), and ovonic unified memory (OUM). Each has impressive backers.

FRAM replaces the DRAM's charge-storage capacitor with a crystal of ferroelectric material. Atoms in the material occupy one of two stable positions (representing one and zero) and change positions in response to an applied electric field. FRAM's backers include **Agilent Technologies** (A), Hitachi, **IBM** (IBM), **Infineon Technologies** (IFX), **Micron** (MU), **Motorola** (MOT), **NEC**, Ramtron, **Samsung**, and **Texas Instruments** (TI). At the last International Electron Devices Meeting (December 2002), Agilent, Ramtron, and Texas Instruments described a 64-Mb FRAM. (For perspective, Samsung described a 2-Gb flash memory at the same conference.)

MRAM replaces the DRAM's charge-storage capacitor with a small magnet. The direction of the magnetic field occupies one of two stable positions, and the field changes direction in response to an applied electric field. MRAM's backers include **Cypress Semiconductor** (CY), Hitachi, IBM, Infineon Technologies, **Mitsubishi**, Motorola, NEC, **Philips Semiconductors**, Samsung Electronics, **STMicroelectronics**, **Toshiba**, **TSMC** (TSM), and **Union Semiconductor**. NEC and Toshiba plan to deliver a 256-Mb MRAM by 2004.

Ovonic memory stores bits in a phase-change material that is similar to the material used in CDs and DVDs. The material's stable crystalline state has low electrical resistance and its amorphous state has high electrical resistance. OUM's backers include **Azalea Microelectronics**, **BAE Systems PLC**, **Intel** (INTC), **Ovonyx**, STMicroelectronics, and Toshiba. Azalea, working with Intel, built a 4-Mb prototype. In February 2003, STMicroelectronics extended its joint-development agreement with Ovonyx. It intends to use OUM in both stand-alone applications and in its microcontrollers.

In addition to FRAM, MRAM, and OUM, there are numerous startups backing novel ideas. Each non-volatile memory candidate has advantages and disadvantages (*Dynamic Silicon*, May 2002). For fifteen years, non-volatile memory candidates have struggled to displace entrenched competition. Until now, it has been a battle they could not win.

Lessons

Flash memory, SRAM, DRAM, and the hard disk have occupied solid positions in the PC for twenty years. Each had advantages that others could not assail. But their dominant market positions derive from the PC's position as the dominant platform. The PC market is shifting from leading-edge PCs to value PCs. This change in profitability for the PC is shifting engineering resources to emerging applications in untethered systems. Untethered systems, the bulk of which are consumer items, balance low cost, battery life, and performance. The design goal in systems is shifting from cost performance to cost-performance-per-watt. None of the current memory and storage components—flash memory, SRAM, DRAM, or the hard disk—is suited to applications that emphasize cost-performance-per-watt. Not surprising for components that evolved for twenty years in a watt-rich environment.

In the mostly tethered PC, these memory and storage components fill complementary niches. Untethered systems want non-volatility and cannot afford the luxury of four component types. There's a huge cost incentive to reduce the set to one working-memory component and one bulk-storage component.

Non-volatile memory candidates have struggled for fifteen years to dislodge PC memory components. The candidates have been unsuccessful because they couldn't beat the incumbents in cost or in density. Aspiring technologies often face the problem of how to close the incumbents' gap. Incumbents have a long head start and they can be rapidly improving. But the value PC and emerging untethered applications change the rules, by shifting the design objective from cost performance to cost-performance-per-watt. Each incumbent has serious shortcomings for emerging untethered applications. The result is enormous incentive to develop FRAM, MRAM, OUM, or some other non-volatile memory, for untethered applications.

There's a similar story for bulk storage. The hard disk occupies an unassailable niche in the tethered PC, and it will continue to do so. But the hard disk is unsuited to untethered applications. The need for bulk storage in untethered devices will underwrite development of MEMS-based storage.

Dense, fast, non-volatile memory will enable rapid growth in untethered applications and *will then displace flash memory, DRAM, and SRAM in the PC.*

I've been talking about how dense, fast, non-volatile memory will change *systems*. It will also have profound consequences in *chips*.

Configuration SRAM limits the convenience and the security of today's programmable logic devices, and it is a limiter of circuit capacity and of performance. Configuration SRAM is inconvenient because SRAM cells don't retain configuration bits across power cycles. Loading configuration bits from a non-volatile source outside the chip requires an extra chip, delays useful work, and reduces security by exposing configuration bits outside the chip. The enormous overhead of four to six SRAM transistors supporting each programmable circuit limits PLD capacity and performance.

The power efficiency and capacity of new non-volatile programmable logic devices will accelerate their application in untethered devices. These advantages will help PLDs displace inherently less efficient microprocessors and digital signal processors. PLDs will be the workhorses in untethered applications.

Untethered devices will usher in a new component set, consisting of non-volatile PLDs, non-volatile memory, and MEMS-based storage.

GILDER TECHNOLOGY REPORT PUBLISHED BY GILDER PUBLISHING, LLC AND FORBES INC. • Copyright ©2003, by Gilder Publishing, LLC

291A MAIN STREET, GREAT BARRINGTON, MA 01230, TEL: (413) 644-2100, FAX: (413) 644-2123 • EMAIL: INFO@GILDERTECH.COM

EDITOR IN CHIEF	EDITORS	EXECUTIVE EDITOR	ANALYSTS	MANAGING EDITOR	ART DIRECTOR	SUBSCRIPTION DIRECTOR
George Gilder	Nick Tredennick Brion Shimamoto	Bret Swanson	Mary Gorski John Hammill	Marie Lavinio	Charles Bork	Rosaline Fernandes

For subscription information telephone toll free: (800) 292-4380 • WWW.GILDERTECH.COM