Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 3 | The Investor's Guide to Breakthrough Micro Devices | March 2002

Microprocessors In Waiting

Every year, the National Football League's players get bigger and faster. The game gets faster and tougher. But the offensive and defensive players change in the same way. Fundamentally, the game stays the same. Imagine a hypothetical NFL, where offensive players got *a lot faster* and a little bigger each year, while defensive players got *a lot bigger* and a little faster. I don't know what would happen, but I know the game couldn't stay the same. Computer systems are like this hypothetical NFL. Progress in semiconductors is opening a rift that changes the game's rules. But this "NFL" wants you to think the game is staying the same.

Improving performance is hard. The transfer rate between Intel's Pentium 4 microprocessor and its memory, for example, is three times what it was for the Pentium III. The latest version of the Pentium 4 has a second-level cache that is twice as large and has three times the transfer rate too. Yet the Pentium 4 at 2 GHz is only about 70% faster than a Pentium III at 1 GHz, despite having twice the clock rate, three times the microprocessor-to-memory transfer rate, twice the L2 cache size, and nearly twice as many transistors (55 million for the P4 versus 28 million for the P3). Diminishing returns.

The vital organs of a computer system are the microprocessor, memory, and hard disk. The microprocessor and the memory improve at a Moore's law rate. Transistors get smaller and faster. But what are these transistors doing? It turns out, often, not much.

Microprocessors run programs, like Microsoft Word or Excel. Word's instructions tell the microprocessor how to manipulate documents. The microprocessor "reads" the instructions and data from memory. Today's leading-edge microprocessors contain tens of millions of transistors organized into processing subsystems. The microprocessor breaks activities into subsystems, called "pipeline stages." Think of them as stations in an automobile assembly line. Chips are getting more complex. The Pentium 4 has twenty pipeline stages, while the Pentium III has ten. The microprocessor's clock coordinates the stages. Each clock tick advances the pipeline's information by one stage. More pipeline stages mean less process-

In This Issue:

The semiconductor world is dividing in two. One part supports tethered devices (ones that need wall sockets for power) and the other part supports untethered devices. In the tethered world, PC performance has peaked. The MHz rating of the microprocessor indicated PC speed—the greater the MHz, the faster the PC. That is no longer true. In the 1980s, semiconductor makers started to optimize their manufacturing processes to achieve different objectives for logic circuits (used in microprocessors) and for memory (DRAM). Logic circuits need to be fast (switch from "1" to "0" or from "0" to "1" at the slightest electrical nudge). Memory circuits need to remember what they are, i.e., they need to be stable—their ones and zeroes have to be coerced into changing, thus taking longer. The result is a huge mismatch between the speed of PC microprocessors and the speed of DRAM, effectively discounting any performance gains brought by speeding up the microprocessor. Cache memories make up this difference in speed, but the gap is getting too wide for them too. The PC microprocessor companies are aware of this but they will not acknowledge it because their pricing is tied to the microprocessor speed expressed in MHz. This situation will continue because of market inertia, but at some point, buyers will realize that the performance benefits are no longer in the MHz. Buyers are better off with, say, a faster graphics card. Embedded microprocessors in untethered devices have always bowed before the throne of energy-efficiency, cost, and just-adequate performance. In untethered devices microprocessor clock frequencies have remained consistent with DRAM access times. There is no mismatch like the one in the PC world. Still, the performance of untethered devices needs to be improved. 3D chips and MEMS promise to speed up storage and to maintain power efficiency. As the need for high-end chips declines, engineering resources will be directed towards untethered devices. Just in time.

ing at each stage, so the clock rate can be higher—but it means that there are more stages to get through to complete an instruction. Carrying out the instruction's orders involves reading, manipulating, and "writing" data. This reading and writing takes time. Because reading happens so much more frequently than writing, we can characterize the computer's behavior by looking at the process of reading. Reading an instruction (or data) has two parts: *finding* it in the memory and *moving* it from the memory to the microprocessor.

	Capacity	Speed
Microprocessors	40%	30%
DRAMs	58%	9%
Hard Disks	58%	12%

Table 1. Annual improvement in computer components.

The time it takes to find it in the memory is called the *access time*. The rate at which instructions or data can be transferred from the memory to the microprocessor is called the *transfer rate*. If your house is like mine, the hot water heater is a long way from the shower head. Access time is the time it takes for hot water to arrive at the shower head. Transfer rate is the number of gallons per hour flowing out of the shower head. Bigger pipes (like wider buses) and more pressure (like more megahertz) improve the transfer rate. Access time is harder to improve. Summary: access time is startup delay; transfer rate is the speed things happen once they start up.

If the instructions and data are in the memory, great; if not, instructions must be transferred from the hard disk to memory first. Instructions and data move from the hard disk to memory as the microprocessor needs them. This takes time since the hard disk also has a characteristic access time and transfer rate.

In the old days—around 1980—memory chips were about the same speed as microprocessors. The original IBM PC, for example, ran at 4.77 MHz, meaning that its internal clock ticked every 210 nanoseconds. The PC's 64-Kb (64-kilobit) memory chips, called DRAM (dynamic random-access memory), had access times of 225 nanoseconds.

By 1984, the PC had a hard disk. The hard disk had an access time of 85 milliseconds. If instructions or data weren't in DRAM, the microprocessor waited about 425,000 clock ticks before the first byte of information moved from the disk to the memory.

Microprocessors and memory

Since the early '80s, microprocessors, DRAMs, and hard disks have gotten bigger and faster—but, like the hypothetical NFL players, they have gotten bigger and faster at different rates. PC microprocessors are now clocked at more than 2 GHz. That corresponds to 400 clock ticks in the space of a single clock tick in the original IBM PC. DRAMs have grown from 64 kb to 256 Mb—a factor of 4,000. The next memory chip, due this year, will be 1 Gb. But, as memory chips get bigger, it gets harder to make them faster (it takes longer to find a book in a million-volume library than it does in a hundred-volume one). Today's DRAM access times are 30 nanoseconds. That is only five times as fast as 1981's chips. In 1981, microprocessors and DRAMs were about the same speed; in 2002, leading-edge microprocessors are sixty times faster than DRAMs.

The difference between the microprocessor's clock rate and the memory's access time is growing. Fig. 1 shows the number of microprocessor clock ticks it takes to access memory. From 1980 through 2001, the number has grown from one to more than fifty.

It doesn't work for the microprocessor to be faster than its memory. The microprocessor spends its time waiting rather than executing instructions. So, as the microprocessor's speed began to pull away from the DRAM, designers added a smaller, faster memory, called a cache, between the microprocessor and the DRAM. Chips for cache memory, called SRAMs (static random-access memories) hold fewer bits, but are

DynamicSilicon

Editors Nick Tredennick, Brion Shimamoto
Publisher Lauryn Franzoni
Designer Julie Ward
Subscription Services Melissa McNally
Editorial Director Richard Vigilante
Chairman George Gilder

Dynamic Silicon is published monthly by Gilder Publishing, LLC. Editorial and business address: 291A Main Street, Great Barrington, MA 01230. Editorial inquiries can be sent to: bozo@gilder.com. **To subscribe call 800-229-2573**, e-mail us at dynamicsilicon@gilder.com, or visit our website at www.dynamicsilicon.com

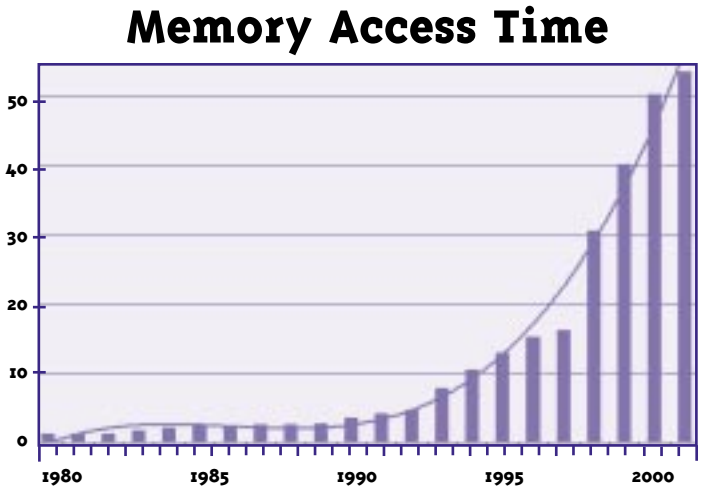

Fig. 1. Between 1980 and 2001, the number of microprocessor clock ticks per memory access has grown from one to more than fifty.

faster than DRAMs.

A DRAM stores information in an array of capacitors. Capacitors hold electrical charge. Each bit consists of one transistor acting as the "gate" to one capacitor. A charged or not-charged capacitor indicates "1" or "0." Since the gate transistor isn't perfect, it leaks. The DRAM chip is "dynamic" in the sense that it loses information unless each bit in its array is periodically read and rewritten; this recharges the capacitor. Semiconductor manufacturing processes for memory create low-leakage, high-voltage transistors, so that the capacitors can be smaller and can carry more charge. DRAM chips are optimized for capacity. Historically, DRAM chips hold four times more every three years.

DRAMs get 9% faster every year. Microprocessors get 30% faster every year. In 1980, microprocessors and DRAMs were about even. As the microprocessor pulled away from the DRAM's speed, it needed faster memory. One way to faster memory is to build the memory with a manufacturing process used for logic circuits. This is SRAM. An SRAM stores information in an array of transistor-based memory cells. Each bit consists of six transistors in a circuit that maintains a "1" or a "0" on its wires. Since SRAMs are built in the same semiconductor process as the microprocessor, they get faster at almost the same rate.

The microprocessor is running away from DRAMs in speed. Double the microprocessor's speed and the PC gets only a little faster; it can't double in speed because the memory holds it back. SRAM chips form a bridge between the fast microprocessor and the slow DRAMs. The tradeoffs in memory chips are capacity, speed, cost, and power. DRAM chips typically have sixteen times the maximum capacity of SRAM chips, but SRAM chips are eight to sixteen times faster. SRAM chips, like the DRAMs, get four times bigger every three years, so, like the DRAMs, their performance suffers as the number of bits grows.

In 1989, a small (8 kB) cache migrated onto Intel's x86 microprocessor chip, speeding the connection between the microprocessor and the cache. The microprocessor was still running away, so designers made the on-chip cache bigger and added a cache for the cache (calling it an L2 cache—for "level-two" cache) off chip. Today's leading-edge PCs have two levels of *on-chip* cache and one L3 cache off chip. The next generation of leading-edge microprocessors will have three levels (cache for the cache for the cache) of on-chip cache.

Several levels of cache match the speed gap between the microprocessor and the DRAM, with smaller, faster caches close to the microprocessor and larger, slower caches closer to the DRAM. But smaller caches hold less, increasing the probability that instructions or data that the microprocessor wants won't be in the cache. If an instruction isn't in the L1 cache (called a cache "miss"), the microprocessor waits for the L2 cache. If it's

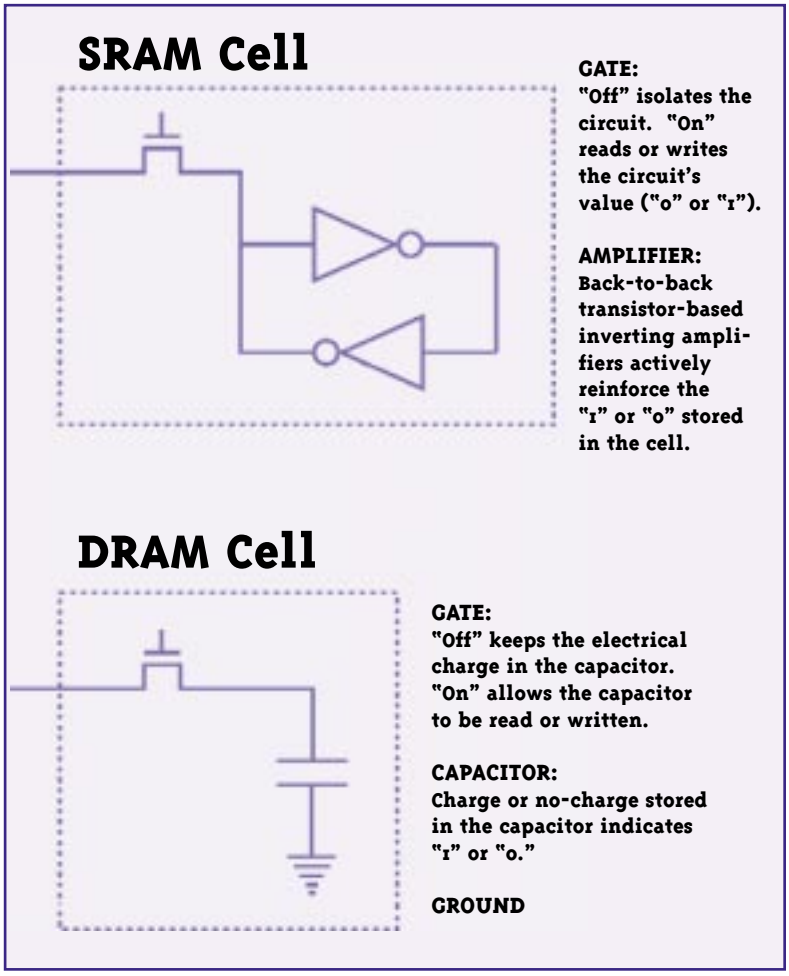

Fig. 2. SRAM cells store a "1" or "0" with transistor amplifiers; DRAM cells store a "1" or "0" on a capacitor.

not in the L2 cache, the microprocessor waits for the L3 cache, and so on. With today's sophisticated microprocessors, this can really slow things down. A 2-GHz, microprocessor capable of starting eight instructions every clock cycle could issue 640 instructions during a single 40-nanosecond DRAM access. In this sense, cache misses are expensive!

Moving DRAM onto the chip with the microprocessor would better match the memory delay to the microprocessor's speed, but it's not a good technology match. The microprocessor needs fast transistors, so it tolerates the current leakage that it trades for speed in the semiconductor process. Transistors that switch fast don't turn completely off. They "leak." DRAM cells are the opposite: they can't tolerate leaky transistors. The semiconductor process for digital logic makes low-threshold-voltage transistors for speed and it tolerates the high leakage currents that result. The process for memory makes high-threshold-voltage, low-leakage transistors.

Corvette enthusiasts are unique among car owners. My Corvette is ten years old; it has wide, sticky tires, 300 horsepower, and a locking differential. It has six speeds. I can smoke the tires and I can exceed any speed limit in any of the top four gears. It has more than enough performance. This year's Z06 Corvette is lighter and faster, has 405 horsepower, and has a raft of performance features. I want one. Chevrolet has trained its Corvette owners. If Chevrolet introduced a Corvette with 2,000 horsepower, I'd want *that*. Never mind that the Corvette I own has more performance than I could ever use. I'm talking about a mentality. More performance is better; if it's offered, buy it.

That's exactly what the PC industry has done to its consumers. In the beginning, the PC didn't have enough performance. Every improvement brought noticeable performance gains. Increase the clock from 8 MHz to a whopping 16 MHz and performance improved by 100%. Old PCs didn't have enough performance, so consumers bought new systems to get the new microprocessor at the faster clock rate. Increase the clock rate and performance jumps again. It still wasn't fast enough. Over time, consumers came to believe that performance is never good enough and to buy faster clock rates. Today, the PC's performance *is* good enough, but consumers have been conditioned to buy faster clock rates, so that's what the manufacturers deliver. The problem is that the microprocessor clock rate is

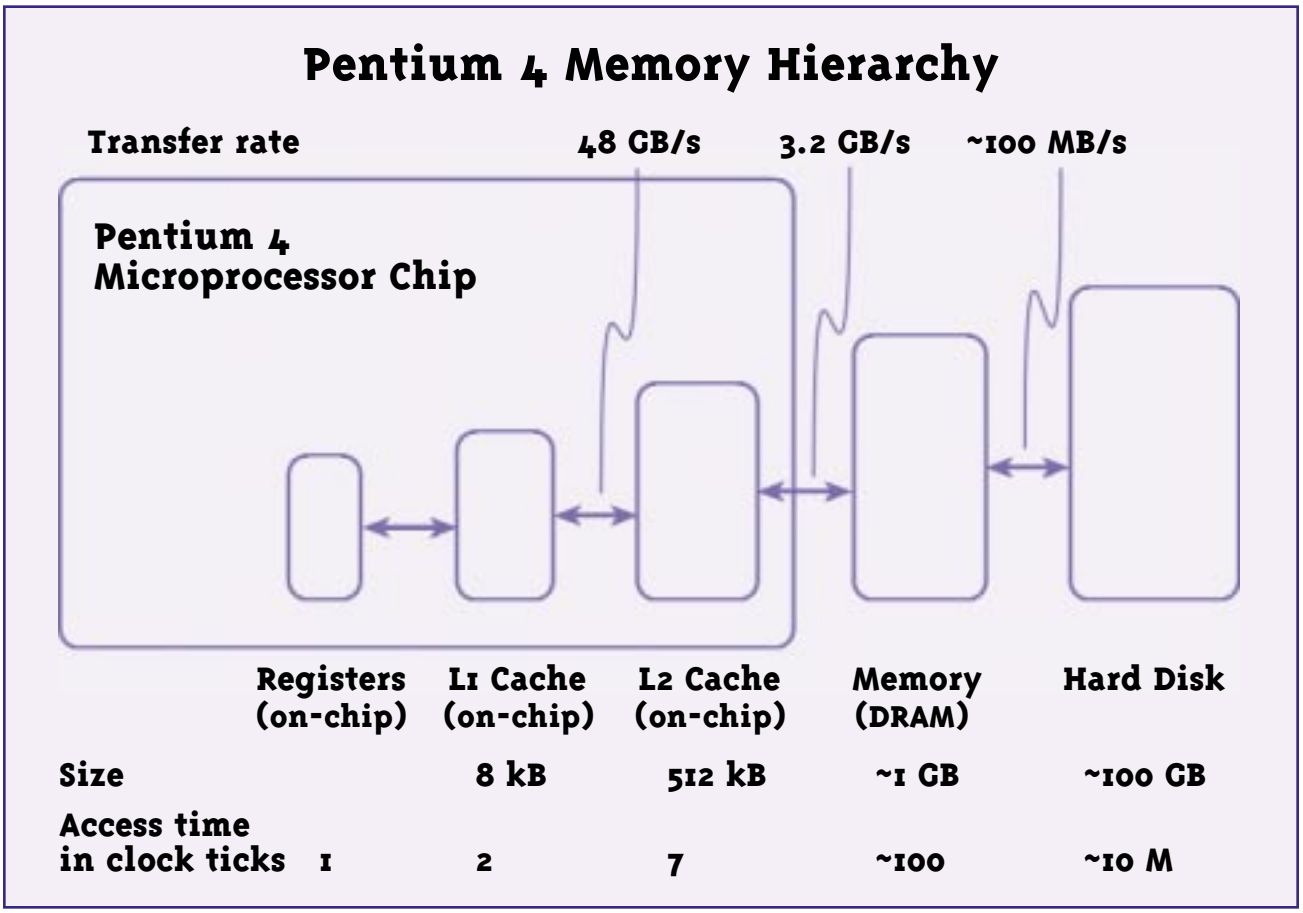

Fig. 3. The Pentium 4's memory hierarchy ranges from one clock tick at the register file to millions of clock ticks to access the hard disk.

no longer an indicator of the PC's performance. The manufacturers aren't about to upset this model; it works for them, so they perpetuate it.

The model is broken. The widening gap between microprocessor clock rates and DRAM access times makes improvements less and less effective. As long as Intel, AMD, Transmeta, and VIA Technologies build microprocessors to optimize circuit speed and the memory makers build memories to optimize capacity, the gap between microprocessor clock rates and DRAM access times will widen. Today's DRAM accesses take about 100 microprocessor clock ticks.

To illustrate the effect of this performance gap, let's assume that just 1% of the misses to the Pentium 4's small first-level cache go to memory (DRAM). The L1 cache access takes 2 clock ticks and the memory access takes 100 clocks. This means there are 99 cache accesses at 2 clock ticks each and 1 memory access at 100 clock ticks, so the microprocessor averages 3 clock ticks per access. *A single DRAM access out of 100 accesses added 50% to average access time.* Now double the microprocessor's speed. The L1 cache access still takes 2 clock ticks (it's on the chip, so it got faster with the microprocessor), but the memory (DRAM) access is now 200 clock ticks (because the microprocessor's clock is twice as fast). The faster microprocessor does 99 cache accesses at 2 clock ticks each and 1 memory access at 200 clock ticks, so it averages about 4 clock ticks per access. The faster microprocessor's 4-clock-tick average is equivalent to a 2-clock-tick average in the slower microprocessor. Doubling the microprocessor's speed raised performance by 50%. For most buyers this is a 50% improvement on something that was already good enough.

If, God forbid, 2% of the microprocessor's accesses have to go to DRAM, doubling the clock speed of the microprocessor raises performance by only 25%. This illustrates the importance of cache design and of branch prediction. Cache design determines the size and the access time of the cache. Microprocessor designers balance the size of each of the cache levels against its access time. A larger cache has fewer misses, but accesses take more clock ticks. Engineers also determine what goes into a cache and, when new information must go into a cache, what will be thrown out. Branch instructions occur as much as 20% of the time. Therefore, branch prediction is critical to keeping the cache filled with the right instructions so that the microprocessor doesn't wait.

Non-volatile storage

As shown in fig. 4, there's a similar story for the hard disk. Since 1984, disk capacity has grown by a factor of

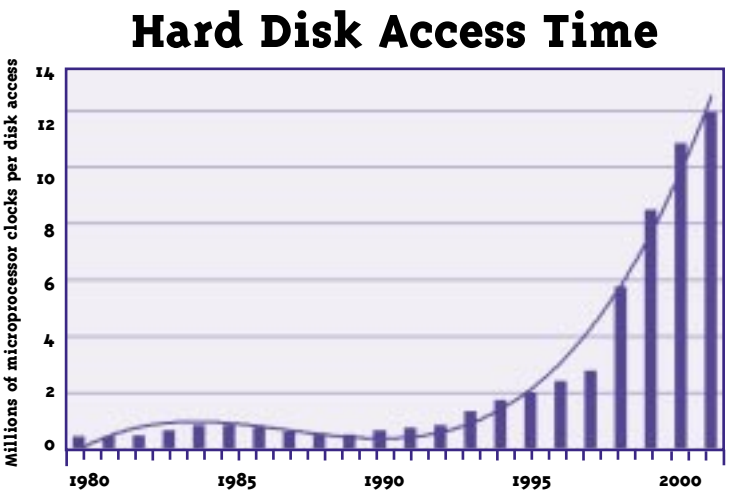

Fig. 4. Between 1984 and 2001, the number of microprocessor clocks per hard disk access has grown from 425,000 to 12,000,000.

10,0000—from 10 MB to 100 GB. The hard disk's access time has improved from 85 milliseconds to 8.5 milliseconds. It's 10,000 times bigger, but only 10 times faster. In 1984, anything the microprocessor needed from the hard disk was 425,000 clock cycles away. Today's wait is 12,000,000 clock cycles.

SRAM and DRAM chips are *volatile*; they lose their information when power is turned off. Hard disks *retain* their information when power is off.

Apple's iPod MP3 player—a leading-edge mobile system—holds a thousand songs. A custom chip containing two ARM microprocessor cores controls its functions and a custom coprocessor on the same chip converts stored music for play. A 5-GB hard disk stores the iPod's songs. It's a Toshiba 1.8-inch hard disk similar to the hard disks found in PC-card inserts for laptops. While it's nice to have a thousand songs, the hard disk isn't ideal for mobile applications. If the disk is spinning, it's using power whether the iPod is playing music or not. If the disk's not spinning, there'll be a delay (seconds) before it can start to play.

Some mobile devices, such as MP3 players, PDAs, and cell phones, can use all the storage they can get. The hard disk isn't great for mobile devices. It is physically too large and it uses too much power. How about solid-state memory? Since semiconductor process improves with Moore's law, won't flash memory eventually be dense enough and cheap enough? Not any time soon. In an 0.13-micron process, flash memory density is 3.8 Gb/sq.in. Hard disks, by contrast, are 30-40 Gb/sq.in. Flash memory is also more expensive at $0.50/MB compared to hard disks at $0.002/MB.

Matrix Semiconductor (*Dynamic Silicon*, Vol. 2, No. 1), using a process for making flat-panel displays, *stacks chips* to improve memory density by a factor of

twelve. Matrix offers 64-MB flash-equivalent chips for $10. At $0.16/MB, that's better, but at 64 MB, it's still far below 5 GB and far above $0.002/MB.

What's needed is storage that's fast, cheap, dense, and energy efficient. Combine the semiconductor's economy of mass production with the density of the hard disk's storage. Enter MEMS-based storage.

I discussed the state of MEMS-based storage some months ago (*Dynamic Silicon*, Vol. 1, No. 5). Not much has changed since then. Nanochip's web site (www.nanochip.com) lists several MEMS-based products. It lists 1-, 2-, and 3-GB MEMS-based storage in a flash-compatible form (those postage-stamp-size cards that fit digital cameras). It lists 250-MB and 1-GB MEMS-based storage chips suitable for embedded designs. These products have specification sheets on Nanochip's web site; the products may be available in the third or fourth quarter of 2002.

Untethered systems

The world is splitting into tethered (plugged into the wall for power) and untethered (dependent on batteries) devices. Tethered devices, with access to unlimited power, can solve the problem of fast microprocessors and slow DRAMs by building a memory hierarchy with SRAMs to bridge the performance gap. That doesn't work for mobile devices that have to run on batteries. They can't afford 2-GHz microprocessors and they can't afford two or three levels of SRAM cache between the microprocessor and the DRAM.

Consumers are the largest market for PCs and for untethered devices (cell phones, GPS receivers, MP3 players). Consumers demand low cost, so both the PC and untethered devices for the consumer market emphasize low cost. PCs sell performance and features; untethered devices sell battery life and features. Embedded microprocessors, the microprocessors for untethered applications, run at a tenth to a hundredth the clock rate of leading-edge PC microprocessors. Embedded microprocessors are slower than PC microprocessors for three reasons. PC microprocessors are fast because that's what their manufacturers have trained customers to buy. Embedded applications can't afford the power, the extra chips, or the cost to build a fancy memory hierarchy that matches a fast microprocessor to a slow memory. And embedded microprocessors only have to have performance that's adequate for the application—for the vast majority of applications, a 30-MHz, 8-bit microprocessor will do.

Microprocessors for embedded applications work either directly with DRAM or they work through small on-chip caches. But cache memory can be a problem in embedded applications. Many embedded applications are real-time systems. A real-time system must meet timing constraints that are set by the world outside the system. It won't do, for example, to have your anti-lock braking system respond in a time that's longer than the minimum for effective action. Application engineers designing real-time systems demand deterministic (fixed, not statistical) behavior from the microprocessor and its memory system. Caches work by holding transient subsets of the main memory's information for fast access. Because what the cache holds depends on the history of system demands, caches introduce non-deterministic behavior to the system and are, therefore, undesirable in real-time embedded systems. The designer of an anti-lock braking system has to know that the microprocessor will respond within, say, 3 milliseconds; it won't do for the system to respond within 2 milliseconds 99% of the time and 50 milliseconds 1% of the time.

High-end embedded microprocessors advertise clock rates of 250 MHz and above, but that is a small application segment.

Embedded microprocessors are in a bind. The software content in untethered systems is rising. SRAMs use too much power and have too little capacity. DRAMs have the capacity, but are too slow. The microprocessor could get faster, but it would burn more power and it would require an expensive, power-hungry memory hierarchy. That leaves making the DRAM faster. One way to make the DRAM faster is to cut the storage array into small pieces and to stack the pieces vertically. Tachyon Semiconductor and Ziptronics do just that (*Dynamic Silicon*, Vol. 2, No. 1). Cutting the storage array into small pieces and stacking them shrinks the distances signals travel. That shrinks the delays that dominate chip access times and it shrinks drive circuits that dominate power use. Stacking pieces improves circuit density, so packaged chip capacity increases. Cutting the chip into pieces also separates the memory array control logic from the memory cells, making both more efficient because they can be built in a no-compromises semiconductor process that's suited to the application. Access time for an experimental 1-Gb synchronous DRAM from Tachyon Semiconductor is 8/4/4/4 (that's 8 ns for the first access and 4 ns each for the next three sequential accesses). That compares with 45/6/6/6 for the best 256 Mb synchronous DRAM competitor (conventional 1-Gb DRAMs aren't available yet).

A host of candidates to replace silicon-based storage are in the research labs. Plastic transistors are one such candidate. All of the potential usurpers that I have looked at, including plastic transistors, have serious

shortcomings. Each candidate has application areas where it shines, such as flexible circuits or special environments, but none has the universal appeal and the wide range of applications to displace silicon in other than niches.

Lessons

Soon the diminishing performance from the widening rift between microprocessor clock rate and DRAM access time will slow the demand for high-end PCs. It'll take a year or so, but the engineering focus will shift from the PC to untethered systems. As it does, the emphasis on high-end microprocessors and on higher microprocessor clock rates will wane. The design emphasis will be on untethered systems characterized by power conservation, by low cost, and by adequate performance (performance above the minimum needed for the application wastes energy). In particular, ascendant applications will benefit Dynamic Silicon microprocessor companies ARM (ARMHY), ARC Cores (ARK), Tensilica, Triscend, and Cypress Microsystems (a subsidiary of Cypress Semiconductor (CY)).

Emphasis on conserving chips, power, and cost favors system-on-chip (SoC) design. SoC design favors the companies, such as Altera (ALTR) and Xilinx (XLNX), that provide prototyping chips. It also favors the chip foundries, such as TSMC (TSM), UMC (UMC), and Chartered (CHRT), as the suppliers both to the programmable logic companies and as the manufacturers of third-party SoC designs.

The shift from tethered to untethered devices strengthens the need for MEMS, to integrate discrete components in the communications circuits and to displace hard disks. MEMS devices for untethered applications are not yet mature, so there are no companies to recommend for RF MEMS components or for MEMS-based storage.

The clock rates for embedded microprocessors will scale with DRAM access times. DRAM performance will improve beyond its historic rate when manufacturers build 3D chips. Tachyon Semiconductor and Ziptronix are developing 3D DRAMs and others, such as Matrix Semiconductor, Irvine Sensors, Mitsubishi, Fujitsu, Kentron Technologies, Amkor Technology, and Tessera (working with Intel), are developing 3D chips for other applications.

Nick Tredennick and Brion Shimamoto
March 15, 2002

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
ARM	Fabless	Excellent	Gains strength as applications move from PCs to untethered devices. Embedded microprocessor cores don't suffer as much from the speed gap as stand-alone chips.
Tensilica	Fabless, Startup	Excellent	Gains strength as applications move from PCs to untethered devices. Soft core and flexible instructions are advantages.
TSMC, UMC	Foundry	Excellent	Move to untethered devices favors designs by fabless companies and portends more business for foundries.
Altera, Xilinx	Fabless	Good	The industry move to system-on-chip (SoC) designs favors PLD companies that provide chips for prototyping and for rapid product introduction. Could be better positioned for portable markets.
ARC	Fabless	Good	Gains strength as applications move from PCs to untethered devices. Soft core and flexible instructions are advantages.
Chartered	Foundry	Good	Move to untethered devices favors designs by fabless companies and portends more business for foundries.
Cypress Microsystems, Triscend	Fabless	Good	Gain strength as applications move from PCs to untethered devices. Microcontrollers don't suffer as much from the speed gap as standalone microprocessors.
Matrix Semiconductor	Fabless, Startup	Good	Positioned well to benefit from chip stacking in Flash-compatible markets.
QuickSilver Technology	Fabless, Startup	OK	Substantial lead in design for portable, adaptive applications, but must get a product to market.
Transmeta, VIA Technologies	Fabless	OK	Increasing demand for midrange systems and for portable systems strengthens markets. Intel's presence in the market makes success more difficult. Both companies should offer x86-compatible cores.
AMD, Intel	Integrated	Struggle	Difficult times as more customers are satisfied with less than leading-edge PC performance.
Apple	Systems	Struggle	Poor strategy by Apple's executives in electing margins over market share in PCs doomed it to niche markets. Apple is too large to grow on products such as iPod.
Nanochip	Fabless, Startup	Struggle	Needs to ship products. Too much time in development forfeits early leadership position.
Tachyon Semiconductor Ziptronix	Fabless, Startup	Struggle	Good technology for DRAM and mixed-process chip stacking, but must overcome semiconductor industry's innovation and licensing barriers.

The "position for the future" and "the way I see it" apply only to the topic of the issue. Possible positions for the future are: excellent, good, OK, struggle, and fail. A company that is "excellent" with respect to horizontal fragmentation of an integrated business may, for example, "struggle" with cultural obstacles in another technical transition. A company listed as "struggle" in another issue could be listed as "good" in this issue since issues cover different topics.

Dynamic Silicon Companies

The world will split into the tethered fibersphere (computing, access ports, data transport, and storage) and the mobile devices that collect and consume data. Dynamic logic and MEMS will emerge as important application enablers to mobile devices and to devices plugged into the power grid. We add to this list those companies whose products best position them for growth in the environment of our projections. We do not consider the financial position of the company in the market. Since dynamic logic and MEMS are just emerging, some companies on this list are startups.

Company (Symbol)	Technology Leadership	Reference Date	Reference Price	2/28/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	19.07	14.66 - 33.60	7.4B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	37.21	29.00 - 53.30	13.6B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.59	4.76 - 6.74	£109M
ARM Limited (ARMHY***)	Microprocessor and System-On-A-Chip Cores	11/26/01	16.59	11.95	8.39 - 19.20	4.1B
Calient (none*)	Photonic Switches	3/31/01				
Celoxica (none*)	DKI Development Suite	5/31/01				
Cepheid, Inc. (CPHD)	MEMS and Microfluidic Technology	12/17/01	4.73	2.68	1.48 - 11.48	71.2M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	22.48	16.06 - 34.0	3.1B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	19.85	13.72 - 28.95	2.4B
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
QuickSilver Technology, Inc. (none*)	Dynamic Logic for Mobile Devices	12/29/00				
SiRF (none*)	Silicon for Wireless RF, GPS	12/29/00				
Taiwan Semiconductor (TSM†)	CMOS Semiconductor Foundry	5/31/01	14.18††	16.25	8.39 - 20.14	54.7B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	3.06	1.17 - 30.38	412M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	8.45	4.25 - 10.41	22.4B
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	11.59	9.71 - 31.5	904.3M
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	35.92	19.52 - 52.14	12.05B

† Also listed on the Taiwan Stock Exchange

†† TSM reported a stock split on 6/29/01. The Reference Price has been adjusted for the split.

* Pre-IPO startup companies.

** ARK is currently traded on the London Stock Exchange

*** ARM is traded on the London Stock Exchange (ARM) and on NASDAQ (ARMHY)

NOTE: This list of Dynamic Silicon companies is not a model portfolio. It is a list of technologies in the Dynamic Silicon 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. The authors and other Gilder Publishing, LLC staff may hold positions in some or all of the companies listed or discussed in the issue.

Ask Nick: Don't forget, all subscribers have exclusive access to Nick on the DS Forum. Just enter the subscriber area of the site and log on with your questions or comments.