Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

Life After Moore's Law

When Moore's law ends, does progress in electronics come to a halt? Will the tech sector look like the aircraft industry after the era of rapid progress in jet-engine design? Is it time to move from tech to the soft-drink sector? My answer is that progress in electronics will continue. Electronics won't need *density* improvements to continue its growth. Electronics-industry growth and Moore's law progress have been Siamese twins because the PC has been demanding all the performance that Moore's law could deliver. As demand for performance in PCs slows, it will become apparent that Moore's law is a *supply law*. The electronics industry will grow through opportunities already available (enabled years ago by Moore's law progress), cutting the industry's dependence on leading-edge progress in semiconductors. In a few years, we'll wonder why we thought Moore's law and electronics-industry growth were so tightly coupled.

Moore's law: the number of transistors on a chip doubles every eighteen months.

Since Moore's law was first stated in 1965 it has been the pace of competition for the semiconductor industry. Our strategy of doubling transistors by making them smaller is hitting fundamental limits. Present day semiconductor fabrication plants ("fabs") coax transistors, capacitors, and wires out of silicon features whose dimensions are *less than the wavelength of visible light*. By 2014, parts of a transistor will be three atoms thick. Transistors can't get much smaller than that. In 2014, the capital cost of a fab will be $50 billion. You can't get much more expensive than that.

A cube of silicon measuring 100 microns on an edge is a *small* grain of sand; its surface area totals 0.06 mm^2. This is less than half the area of the period at the end of this sentence. In today's leading-edge, 130-nanometer (0.13 microns) semiconductor process, that's enough room for *50,000 transistors*. If we follow the industry's forecast of 22 nanometers (0.022 microns) in 2016, *two million transistors* fit on that grain of sand. Two-million-transistor grains of sand are beyond practical utility for the vast majority of applications. Our world scales from dust mites to elephants. We don't notice dust mites, but we pay attention to scorpions. Tiny transistors are as unnoticeable as dust mites—they only interact with things of comparable size. Tiny transistors have to scale up to interact with us.

In 2002, the leading-edge logic process will be 130 nanometers (0.13 microns). Since chips are being manufactured in this process, last year's issues with 130 nanometers must have been solved. The next-generation logic process, 100 nanometers (0.1 microns), will go into production in 2003. Concerns remain about how to manufacture these chips, but there are few issues and the solutions, while difficult, are known. Beginning in 2005, with the 80-nanometer generation (that's 0.08 microns, *one-fifth* the wavelength of visible light), there are issues for which no solutions are known. Each subsequent generation has more issues with no known solutions. The industry's name for the collection of issues with no known solutions is "the red-brick wall." This wall has always receded with time.

Life after CMOS

There's concern over what happens to computing at the end of Moore's law and about where the industry will go after CMOS. This concern drives research into new computing methods and into new devices (things that act like transistors). Industry, government, and academia are researching quantum devices,

Vol. 2, No. 1
January 2002

molecular devices, DNA computing, biological computing, amorphous computing, carbon nanotubes, single-electron transistors, plastic transistors, and a host of others. Is there something on the horizon that will sneak up and displace silicon and CMOS? All of these methods and devices have advocates and advantages; all also have major drawbacks. Quantum computing could be orders of magnitude better at solving cryptography problems than conventional computers, but its advantages in other applications are unknown or unimpressive and most of its implementations will require bulky, expensive cryogenic cooling. Molecular devices, which take one of two shapes based on applied voltage, promise non-volatile storage. Molecules offer huge increases in storage density and in energy efficiency, but research is a long way from practical devices. There's a similar story for each alternative: huge gains in some applications, impressive characteristics for some device attributes, and major shortcomings in several areas.

Some alternatives, such as molecular storage or plastic transistors, may carve niche applications, *but nothing I've seen has the potential to replace or to overtake silicon.* The incumbent silicon process still has enormous potential. Intel has demonstrated a silicon transistor with a 15-nanometer (0.015 microns) gate operating at 2.63 *terahertz* (2,630 gigahertz!). Besides its impressive speed, the transistor reduces off-state current leakage by a factor of 100 and it reduces gate leakage by a factor of 10,000. These impressive improvements in leakage currents and in frequency of operation breathe new longevity into silicon applications for low-power and for high-frequency applications (leakage currents generate heat and they dominate power loss in many low-power applications). Intel plans to manufacture chips with some of these characteristics as early as 2005.

If "computing" means programs running on supercomputers or even on personal computers, then the answer is "who cares?" because it's inconsequential in unit volume. If "computing" means programs running on embedded microprocessors, then the transition from CMOS will be at least ten years after Moore's law ends.

Any transition to non-CMOS processes will be smooth. Electronic systems will co-opt solutions from nature. Nanomagnetics, for example, uses nature's iron-storage protein (the ferritin molecule) for magnetic-storage media by replacing the iron with a cobalt-platinum alloy. A hexagonal-close-pack of modified ferritin proteins could achieve 45,000 Gb/sq.in., well above today's ~40 Gb/sq.in. limit. Think of a portable MP3 player holding a million songs. Designers will borrow efficient molecular and biological mechanisms from nature, but the transition will be gradual.

Supply and demand

Fig. 1 shows a hypothetical twelve-year window in the PC market. The "supply" line shows the PC's performance doubling every two years. Corresponding to the supply curve is a difficult-to-see and often-forgotten *demand* curve that represents the user's performance requirement (*Dynamic Silicon*, Vol. I, No. 3). In this hypothetical example, the PC's performance begins at unity and the user's performance requirement begins at ten. The users' expectations for performance rise with time as users, operating systems, and applications become more sophisticated. There is, however, *no necessary connection* between the rate of increase in expectations (demand) and the rate of increase in performance (supply).

If the demand for performance begins above the performance being supplied and if it grows faster, then users will always demand more performance than the industry can deliver. If the demand for performance grows more slowly than the supply of performance, then the lines will cross and performance will exceed demand. Users will then no longer pay a premium for additional performance, decreasing the incentive to supply more performance.

The demand curve isn't the same for all users; some need all the performance a leading-edge PC can provide, others don't. Early PC adopters have high expectations and rapidly rising expectations; late adopters have lower expectations and slowly rising expectations. PC users fill the area between the "Rapidly Growing

DynamicSilicon

Editors	**Nick Tredennick**
	Brion Shimamoto
Publisher	**Lauryn Franzoni**
Web Editor	**Jorin Hawley**
Designer	**Julie Ward**
Subscription Services	**Melissa McNally**
Editorial Director	**Richard Vigilante**
President	**Mark T. Ziebarth**
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

Demand" and "Slowly Growing Demand" lines of fig. 1. Further, as the user base expands it brings in users with lower performance needs (late adopters)—broadening the range of performance demand and slowing its growth rate.

PC performance improves at a rate close to the Moore's law improvement in the PC's semiconductor components. In the PC's early years, users wanted so much more performance than the PC could deliver that

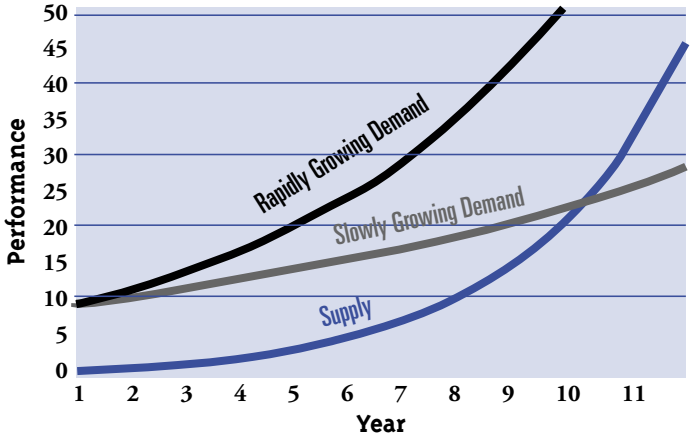

Fig. 1. PC performance (Supply) rises at a rate close to the Moore's law improvement in semiconductors. User expectations (Demand) rise with time, but vary among users.

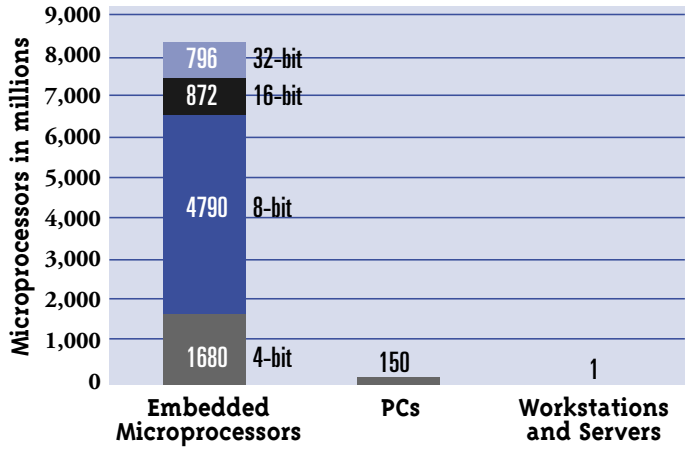

Fig. 2. Embedded microprocessors dominated microprocessor shipments in 2000.

PC makers assumed the demand for performance was infinite. But the demand for performance isn't infinite; demand spans a range. Over time, PC performance improves *and* consumer expectations rise. The PC's performance rises faster than consumers' expectations. So, over time, the PC's performance satisfies more and more consumers. "Value" PCs, which are based on microprocessors a generation or two behind the leading edge, *increase* their share of new computer sales against leading-edge PCs.

The advance of semiconductors is similar to the PC market. Moore's law supplies more transistors per chip or it supplies more performance. And, like the PC industry, the semiconductor industry has been built on an assumption of infinite demand. The advance of semiconductors is so similar to the PC market that we can recycle fig. 1. Moore's law advances in semiconductor process supply performance. Applications (PCs, engine controllers, electronic games, electronic toys, etc.) span the range from slowly growing demand to rapidly growing demand. The PC's microprocessor and memory have defined the line of rapidly growing demand since the PC's introduction. Because the PC business consumes half of the world's semiconductor output it's easy to mistake the advance of Moore's law with the growth of the electronics market.

Fig. 2 shows estimated shipments of all microprocessors in 2000. The first bar is the stack of 4-bit, 8-bit, 16-bit, and 32-bit embedded microprocessors. The second (barely visible) bar is the 150 million PC microprocessors. The third (invisible) bar is approximately one million workstations and servers. Workstations and servers consume high-end, high-margin microprocessors. PCs consume high-volume, high-margin microprocessors. Embedded applications consume high-volume, low-margin microprocessors. The average selling price for a PC microprocessor is about $200. Average selling prices for embedded microprocessors may be less than a dollar and vary substantially depending on features, but probably average less than $4. Embedded microprocessors dwarf PC microprocessors in unit volumes, but the PC microprocessor's high average selling price makes the dollar value of its market about the same size as the market for embedded microprocessors.

Fig. 3 shows important segments of the electronic systems market, classified by design objective. Zero cost is the consumer segment. It is cost sensitive and it is by far the largest segment. In the zero-cost segment, the designer's objective is zero cost of materials. In the zero-power segment, the design goal is zero power dissipation. It includes battery-powered items such as watches and smoke detectors. Most of the zero-power segment is inside the zero-cost segment. Designs in the zero-delay segment strive for fast response (zero delay from request to answer). The zero-delay segment includes personal computers and GPS receivers. Most of the zero-delay segment is inside the zero-cost segment. Personal computers, for example, are inside the zero-delay segment and are *inside* the zero-cost segment. Workstations and servers are inside the zero-delay segment and are *outside* the zero-cost segment.

Since the PC is a consumer product that competes on performance, it is in the overlap of the zero-cost segment with the zero-delay segment. The PC's microprocessor and its memory chips hug the leading-edge of the demand curve that corresponds to the Moore's law supply curve. Performance *and* cost matter. The growth of PC- and memory-based applications depends on Moore's law advances. For most of the products leading the growth of the zero-cost segment, however, only cost matters. Trailing-edge microprocessors deliver all the performance that's needed. The growth of zero-cost applications does not depend on further Moore's law advances. Past advances in semiconductor processing *enabled* those applications years ago.

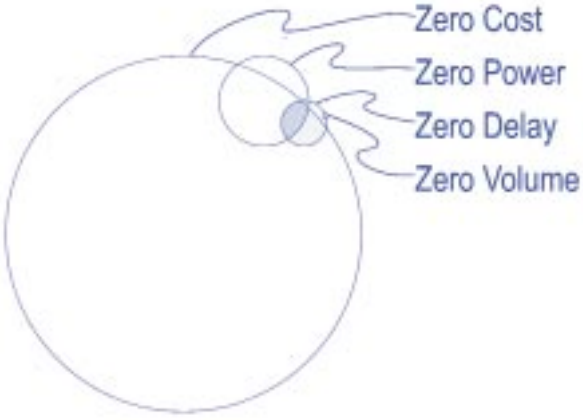

Fig. 3. The electronic systems market is dominated by the consumer segment with a design goal of zero cost. PCs are in the overlap between zero cost and zero delay.

The year the PC was introduced, 1981, it sold about 15,000 units. That same year, manufacturers shipped fewer than 450,000 microprocessors (my estimate). By 2000, the market for PCs had grown to 132 million units a year, showing a compound annual growth rate of just over 60% per year for the period. Over the same period, the microprocessor market grew to 8.5 billion units a year, showing a compound annual growth rate just under 70%. The market (in units) for embedded microprocessors grew even faster than the personal computer market even though microprocessors started from thirty times the base! The importance of this is that the electronics industry has slaved itself to the PC, improving the semiconductor process at a Moore's law rate to build the personal computer market. The high end of the PC market *defines* rapidly growing demand for semiconductor performance. As the semiconductor industry races to satisfy this demand, it leaves a performance wake that enables all of the applications between the limits of the demand curves.

Fig. 4 shows the range of engineering problems. The x-axis is the problem size from small to large. The y-axis is the required solution speed (a guided missile has to solve its navigation problems faster than it flies). Moore's law improvements in semiconductor processing expand the range of economical solutions (as indicated by the arrows in fig. 4). At the high end, microprocessors get faster and they become more capable, so leading-edge microprocessors solve harder problems. At the low end, microprocessors get cheaper, enabling more low-end applications. Moore's law advances in semiconductor process enable a wider range of problem solutions. But the area enabled for microprocessor applications represents the range of *potential* applications, not the applications themselves. Implementation requires engineers. The area of economical solutions enabled by Moore's law progress is *running well ahead of engineers' ability to exploit it*.

The leading edge and the trailing edge

The press, research organizations, and high-visibility industry segments, such as PCs and workstations, focus on the leading edge. The belief is that without advances at the leading edge, the industry will cease to grow—thus concern over an end of Moore's law.

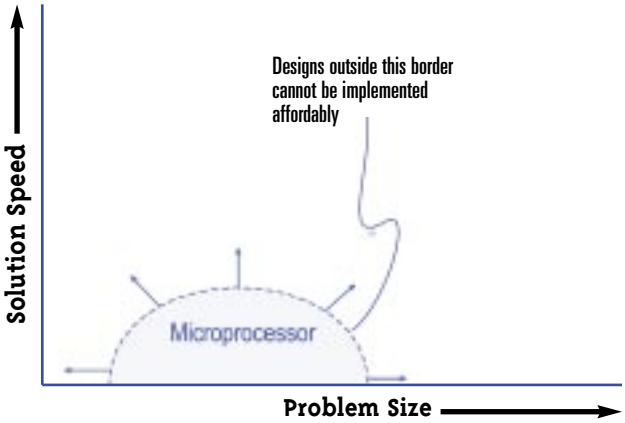

Fig. 4. The microprocessor is an economical solution for a range of engineering problems.

Concern may be misplaced; the leading edge isn't the industry. Leading-edge computers are zero percent of unit volume. Personal computers consume less than two percent of the microprocessors that ship each year. Each year, manufacturers ship more than one microprocessor for every living person on the planet (six billion people, eight billion microprocessors). Most of the eight billion microprocessors are "embedded," meaning that the microprocessor is hidden inside a system and is invisible to the user. These aren't Pentium 4s, or Itaniums, or

even ARM 9s, they're ancient 8051s, 6805s, and Z80s. The creaking-old 8-bit microprocessors that make up the bulk of computing that's embedded in systems each year aren't based on leading-edge processes, they're trailing-edge, pad-limited designs cranked out in 0.5-micron (500 nanometers!), fully depreciated fabs.

Leading-edge semiconductor process in 2002 is 130 nanometers (0.13 microns). At Chartered (CHRT), the third-largest foundry, only 4% of revenues come from semiconductors at 180 nanometers (0.18 microns) or better. At TSMC (TSMC), the world's largest foundry, demand for leading-edge process accounted for only 5% of its wafer starts in 2000. In 2004, when the industry's leading-edge process will be 90 nanometers, TSMC expects demand for the three most recent process generations (130, 107, & 90 nanometers) to account for only 20% of its wafer starts.

Applications for low-end embedded microprocessors continue to grow and will continue to grow—they trail leading-edge processors and processes by three generations or more. Manufacturers of the most popular twenty-year-old microprocessors still ship hundreds of millions of units per year. *An end to Moore's law won't be felt in the majority of embedded applications until ten years after it happens* (I argue below that it will never be felt). The market for these low-end microprocessors is representative of the invasion of electronics into new systems (toothbrushes, hair dryers, transmissions, bumpers, etc.). The rate of this invasion is limited by the availability of design engineers, not by the lack of leading-edge microprocessors.

Size for minimum utility

Circuits and the real world. Spoons and shovels are similar devices on different scales. Scale is important to utility; it wouldn't be easy to eat with a shovel or to dig trenches with a spoon. Scaling is important to utility in electronics as well. The vacuum-tube-based walkie-talkie of World War II was bulky and unreliable. Transistors shrunk the mobile radio, improving its reliability and its utility. Integrated circuits continue this process. By 2016, we could put two million transistors on a small grain of sand.

Do the smaller transistors of each generation of chips improve reliability and utility? Let's say you could shrink the transistor to a single-electron device. Is that better than a transistor that operates with a few thousand electrons? Not necessarily. Just as a miniature spoon wouldn't work well for soup, a transistor that's too small lacks the ability to interact with the real world. It wouldn't do to shrink the cell phone's electronics onto a grain of sand. The cell phone's transmitter has to put out enough power to reach the base station and its speaker has to generate sounds you can hear in a crowded airport.

Sensors and actuators and the real world. Sensors and actuators measure and interact with the real world. If they're too big, they use too much power and they waste resources: if they're too small, they are inaccurate or ineffective.

Laboratory chemical analysis requires test-tube-size samples and similar amounts of reagents, and it creates test-tube-size wastes. Chip-based chemical analysis works with samples a thousand times smaller, uses correspondingly smaller quantities of reagents, and creates (and confines) smaller wastes. Chip-based chemical analysis can be faster, cheaper, more accurate, more automated, and safer than laboratory analysis. Smaller is better—but only to a minimum that retains utility. Shrink the sample size too much and the sample won't be representative.

Microelectromechanical sensors and actuators interact with the real world as the collectors and consumers of data. These sensors and actuators have a minimum effective size that's within the capability of today's fabrication processes. Microprocessors and other electronic components that interact with the sensors and actuators need power and size that enable efficient interaction. While there are applications for extreme computing and for extreme power conservation, transistors that interact with the environment and with the user will be larger than the smallest that can be made by following Moore's law to its limit.

3D

In the 1954 science-fiction movie *Them!*, eight- to twenty-foot ants, the mutant result of nuclear testing in a nearby desert, invade Los Angeles. It wasn't the first giant-creature movie. The theme has been repeated countless times since. Science geeks, bent on spoiling our entertainment, point out that these giants aren't viable in the same proportions as their lesser real-world relatives. Body mass, they say, increases with volume, but skeletal structures support weight on cross sections. For this reason, the leg bones of an elephant are fatter than those of a scaled-up horse would be. The legs of a twenty-foot ant would be too fragile to support the ant's weight.

Chips and transistors are similar to the scaling of skeletons and mass. We lay out chips in two dimensions. Soon, wires and signaling delays dominate the chip. Process improvements shrink the transistor and allow the maximum chip size to increase. Results combine unfavorably; smaller transistors drive signals on longer wires. Transistors, like marbles, pack better in three dimensions than in two. It's time to go vertical.

Companies like IBM, Irvine Sensors, and Dense Pak build three-dimensional (3D) integrated circuits by stacking chips. It hasn't been popular for other than special applications. There are good arguments against 3D chips. Power dissipation, which is already a problem with two-dimensional chips, is worse in 3D. Yield is the product of yield for each individual chip times the yield for building the assembly. Stacking four chips with 90% yield gives only 66% yield, with perfect assembly. Testing is problematic. Interconnections between stacked chips have been limited to the perimeter.

Changes in the semiconductor process offer new opportunities.

The drive to increase the number of two-dimensional connections among on-chip circuits drove the need for more metal layers. Today's semiconductor processes support eight wiring levels. That's already the third dimension, but it's wires, not transistors. Stacking multiple levels of wires across the chip began to get difficult. An insulating layer covers the first level of wiring. Another insulating layer covers the second level of wiring. The surface of the chip gets bumpy in the third dimension as wires follow hills and valleys created by wires beneath them. Making connections up and down the chip's contours became difficult. IBM invented chemical-mechanical polishing to solve the problem. This polishes each layer flat before the next layer is added. The original purpose was to keep the wafer's surface flat so that the mask illumination would remain in focus everywhere. The result is perfectly flat chips, suitable for stacking.

3D chips on the way

Fabless semiconductor startup Matrix Semiconductor (www.matrixsemi.com) makes 3D chips. Matrix, takes advantage of chemical-mechanical polishing; it also exploits advances in flat-panel displays. Manufacturers of flat-panel displays developed methods for building millions of thin-film transistors on glass. Matrix substitutes a silicon wafer for the glass and stacks multiple layers of thin-film transistors. The company has made chips with transistors stacked to twelve layers and believes that sixteen-layer chips are feasible.

Its first product will be a 64-MB "write-once" memory in a package that is interchangeable with standard flash memory cards. Its price will be about $10, well below comparable flash memory at $38. Since the circuits are stacked, the chips are smaller and cheaper than their 2D counterparts. Each wafer produces ten times as many chips as a 2D layout. The process is close enough to standard that the wafers can be made by foundry partner TSMC on its 0.25-micron process. The company believes that its technology can be extended to general-purpose logic, but its best fit is with write-once memory.

Memory chips are a design challenge because they combine memory (memory cells) with logic (drivers, receivers, and control logic) for the memory array. Memory cells need low current leakage and high (switching) voltage; logic needs low (switching) voltage and tolerates high current leakage in exchange for faster switching. Control logic for the memory array is typically 50% of the chip. One reason it is so large is that the drivers and receivers for the memory array drive wires across the entire 2D array. Private company Tachyon Semiconductor (www.tachyonsemi.com), with patents licensed from Glenn Leedy's Elm Technology, developed a stacked memory with significant benefits over 2D memory chips. Here's the process.

Build the memory controller on a separate chip with a logic manufacturing process. Divide the memory array into eighths (or some other fraction), matching these small memory arrays to the (reduced) size of the controller chip. The logic wafer serves as the base; a memory wafer is flipped upside down and bonded to the logic wafer. This memory wafer is thinned to about 10 microns and another memory wafer is flipped and bonded onto the stack. Memory wafers can be bonded and stacked to any reasonable height (Tachyon's drivers and receivers can drive up to thirty-two memory layers). Since even 2D wafers are thinned before packaging, a typical 3D chip is no thicker than a standard 2D chip. With the 3D memory chip, distances to the memory cells are reduced over what they would be in a 2D layout. Instead of driving across a large 2D chip, they drive a few microns vertically and then across a fraction of 2D surface. These reduced distances speed the chip and reduce power dissipation (the drivers and receivers can be smaller). The memory chip has higher yield since it is smaller. Metal wires connect between layers through vertical channels. There's room for hundreds of thousands of vertical connecting wires.

Tachyon announced a 1-Gb synchronous DRAM chip in November (*Electronic Design*, 19 November 2001) with initial access times more than five times faster than access times for comparable 2D memory chips (8 nanoseconds vs. 45). Subsequent accesses are 50% faster. These chips may be delayed getting to market (due to licensing and funding issues), but the concept is sound.

Stacking thirty-two wafers sounds like a good way to reduce the yield to zero. Yield doesn't go to zero. The chips employ redundancy and self-test. Even the controller need not be a weak link. If extreme reliability is

needed, stack two controllers and a couple of extra layers of memory. The chip can self-test and configure itself on startup. This 3D stacking could mix analog, digital, and memory layers, with each wafer built in a manufacturing process optimized for its function.

3D stacking would enable Altera (ALTR) or Xilinx (XLNX) to leap *above* Moore's law progress in density with its programmable logic devices or perhaps to stack configuration layers on the logic. It would also enable microprocessor manufacturers such as Transmeta (TMTA), Intel, or AMD, who now build families of chips, to build a single processor and then offer a family of chips with multiple processors, single processors, a range of cache sizes, and even on-chip memory. Moving the memory into the stack with the processor removes the major performance bottleneck and it reduces power consumption.

Ziptronix (www.ziptronix.com), a North Carolina startup, has a process similar to that developed by Tachyon. Xilinx is an investor. Both Tachyon and Ziptronix use existing semiconductor processes such as those available from foundries such as TSMC and UMC (UMC).

Electronic Design named both Tachyon and Matrix Semiconductor among its top ten companies to watch in digital integrated circuits (*Electronic Design*, 7 January 2002). Funding, industry inertia, and licensing issues may delay its appearance, but as a technology, 3D is ready to go.

Lessons

There's gloom-and-doom talk about the end of Moore's law. Some of the talk comes from researchers who want money to investigate new devices or new computing models. Some of the talk is born in "the red-brick wall" of the semiconductor industry's forecast. The semiconductor industry has a history of pushing back the red-brick wall. Companies on the Moore's law treadmill push the wall back with investment. Intel's "TeraHertz" transistor shows that the wall can be pushed past 2005. Some of the talk comes from a belief that an end to Moore's law improvements will stall the growth of the electronics industry.

Silicon will dominate the chip industry to beyond the horizon of the industry's fifteen-year forecast (2016). It won't be displaced and the electronics industry will continue its expansion even if Moore's law stalls. That is so for two reasons. The first is supply and demand. The second is size for minimum utility.

Moore's law defines the supply curve for the electronics industry. Moore's law progress supplies performance or it supplies transistors per chip. Moore's law doesn't *drive* growth in electronics; it *enables* growth. The PC market, with its own supply and demand curves, has been driving demand for more performance. As more PC users are satisfied with less than leading-edge performance, demand for leading-edge development weakens. Companies concentrated on the PC market because that's where the money was. Because the PC's microprocessors and memory have been at the leading edge, *Moore's law improvements have left a huge wake of enabled but unexploited areas for growth in electronics.* As the PC market slows engineers move to wireless devices, or to network systems, or to smart appliances, or to interactive toys, and so on. The bottleneck in growth of the electronics industry isn't Moore's law; the bottleneck is limited engineering design resource.

In a cell phone of 250 components, only 12 are integrated circuits. The rest of the components are analog and RF discrete components. There's more to be saved in power and efficiency by working on the discrete components than there is in shrinking circuits inside the 12. Even if all the digital logic fits on a grain of sand, the electronics must have the energy and scale to interact with the real world (drivers for the display, the antenna, and the speaker and sensors for location, for keys, and for the microphone).

Moore's law increases performance and it increases transistors per chip by scaling in two dimensions. 3D circuits offer similar benefits in scaling. It should be OK that it isn't coming from shrinking the chip's 2D features. As hardware gets softer (*Dynamic Silicon*, Vol. I, No. 6), functions migrate to regular arrays such as memory and programmable logic devices. Regular arrays benefit most from 3D implementation, so Altera and Xilinx should benefit from developments by companies such as Tachyon and Ziptronix. Chip stacking leads to fewer unique chips as customization moves from fabrication to custom stacking. Regular arrays and fewer chip types benefit foundries such as TSMC, UMC, and Chartered, which get larger production runs between line configurations.

Moore's law is a supply law. It supplies more performance or more transistors and, at its limit, it will have supplied enough performance and enough transistors. Moore's law has already left in its wake vast areas of enabled but unexploited opportunity. As the PC market matures, the industry's design engineers will shift to systems with embedded microprocessors. Design emphasis will shift from performance to capability, versatility, and portable utility.

Nick Tredennick and Brion Shimamoto
January 21, 2002

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	12/31/01 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	21.22	14.66 - 34.69	8.2B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	44.39	29.00 - 64.00	16.1B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.46	£0.25 - 3.27	£128M
ARM Limited (ARMHY***)	Microprocessor and System-On-A-Chip Cores	11/26/01	16.59	15.59	8.39 - 26.13	5.3B
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	4.20	1.48 - 11.48	111.6M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	26.44	16.06 - 37.13	3.6B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	19.93	13.72 - 29.25	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††	17.17	8.39- 19.49	57.8B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	2.29	1.17 - 37.25	308M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	9.60	4.25 - 10.86	25.5B
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	17.91	9.71 - 38.75	1.4B
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	39.05	19.52 - 59.25	13.0B

† 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.