GILDER
TECHNOLOGY
REPORT

April 2003 / Vol. IX No. 4

# Sea Change In Semiconductors

**The next phase: smaller does not mean better.**

A sea change looms that is as significant as the integrated circuit. The change is caused by the appearance of the value PC and by the appearance of the value transistor. These value plateaus signal declining demand for leading-edge PCs and for smaller transistors. By the metric of cost performance, the personal computer has become "good enough" and the transistor has become good enough. This is not the end of electronic progress. Forty-some years after the first integrated circuit, we are entering a new phase.

PC advances and transistor advances are no longer must-haves for everyone; leading-edge features attract diminishing numbers of customers. The PC, an essentially tethered-to-wall-power device, is ceding its position as the premier platform of electronics to untethered consumer devices. The goodness metric is shifting from cost performance to cost-performance-per-watt. Our familiar electronic components—microprocessors, digital signal processors, today's programmable logic devices, flash memory, SRAM, DRAM, hard disks—were nurtured in the PC's watt-rich environment and score poorly against the new metric. This creates opportunities enabled by new non-volatile memories and by a novel manufacturing process that takes Moore's law into the third dimension.

Today's memory and storage components are unsuitable for untethered devices. Microprocessors and digital signal processors are unsuitable. And today's programmable logic devices are unsuitable. How did we come to this juncture? Where do we go from here?

(This report links themes from recent issues: November and December 2002 of *Dynamic Silicon* and February and March 2003 of the *Gilder Technology Report*. Because the themes' importance is cumulative, I'm devoting all the space in this issue to their combined story. For the April *Telecosm Technologies* information, see www.gildertech.com.)

## Moore's law and the integrated circuit

Moore's law was the supply-side business model. The integrated circuit was the breakthrough. Together, they enabled the rapid growth of the semiconductor industry.

Moore's "law" isn't a law; it's the rate *the industry sets* for itself to shrink transistors. Historically, this meant doubling the number of transistors per chip every eighteen months. Since silicon area, not transistors, determines cost, Moore's law has meant more transistors for the same cost, or the same number of transistors at lower cost. Moore's law answered the question: "How do you differentiate the next widget?" No one had to worry about how to advance product capability because Moore's law did it for them. Since capability is determined by transistors, Moore's law is strongly associated with making products cheaper and better. That's the theory (more later on theory versus practice).

Instead of designing with individual transistors, engineers designed with macro-circuits comprising hundreds or thousands of transistors. Engineers' productivity rose because designing with such integrated-circuit "macros" was more

**Inside:**

efficient than designing with individual transistors. So, integrated-circuit macros proliferated. And with Moore's law, integrated-circuit macros grew in variety and in complexity.

## The microprocessor

Integrated-circuit macros evolved into the microprocessor. The microprocessor was not invented to be the heart of a *computer* system; it was invented to consolidate what had been many separate integrated-circuit macros onto one chip. (Computers were still rare; microprocessors were used in embedded—non-computer—applications.) In the form of a microprocessor, these formerly separate integrated circuits were selectively invoked using a control idea from computer systems: instructions. Thus, the microprocessor brought programmability to *embedded* systems. Instead of building functions from integrated circuits, engineers wrote programs—sequences of instructions—for microprocessors.

A microprocessor, memory, and some peripheral chips consolidated and displaced collections of integrated-circuit macros. Microprocessors thus *mimicked* hardware circuits for a wide range of applications. Granted, the way microprocessors did this was inefficient, serially invoking one integrated-circuit function at a time and using lots of power to do it. But they worked well enough and they were flexible—engineers could make last-minute changes by changing instructions (bits in memory) instead of having to change physical circuits. As with integrated-circuit macros, engineers' productivity rose.

After serving ten years in embedded applications, the microprocessor became powerful enough to be the central processing unit of a computer system. The IBM Personal Computer debuted in 1981. IBM did not invent the personal computer; more importantly, IBM legitimized the idea of personal computing. The personal computer split microprocessor design: some microprocessor makers served the cost-sensitive embedded-systems market and some began designing microprocessors for the performance-oriented computer market.

Microprocessors for the computer market dominated computer research because computer scientists wanted to work on leading-edge designs. These microprocessors monopolized press coverage because the few companies building leading-edge microprocessors commanded high profit margins.

Microprocessors and digital signal processors were developed in a cost-performance environment. Their strong suit is flexibility. To achieve that, they sacrificed computational efficiency and energy efficiency at a time when those traits mattered little. Both are so computationally inefficient and so energy inefficient that they cannot be tweaked to meet the cost-performance-per-watt requirements of untethered systems.

## PC story

The PC's evolution bred three memory components (flash memory, SRAM, and DRAM) and one storage component (the small hard disk). Each component occupies a niche in the PC's system structure. Each has competencies that secure its niche and shortcomings that confine it to its niche.

Flash memory's competence is its non-volatility; it retains stored values even when the power is off. Flash memory holds special programs that initialize the chips on the PC's system board. Though flash memory is slow and wears out over time, it finds good use in the PC because it is used only when the PC restarts.

Dynamic random-access memory (DRAM) is the working memory for PC programs. DRAM's competence is its capacity—as working memory. DRAM is volatile (loses its values on power off) and it is slow. In the original PC, the DRAM and the microprocessor were about the same speed. (Neither one waited for the other.) Over time, companies designed DRAM for higher capacity and they designed the microprocessor for greater speed. Optimizing capacity versus speed led to the speed mismatch between the working memory and the microprocessor (*Dynamic Silicon*, March 2002).

Enter static random-access memory (SRAM). SRAM sits between the DRAM and the microprocessor to bridge the speed gap. SRAM is made of the same transistors that comprise the microprocessor, so it's as fast as the microprocessor, but it has only one-sixteenth the capacity of DRAM. SRAM is expensive, it's volatile, and it's not as dense as DRAM. It takes six transistors to hold a bit in SRAM; one transistor holds a bit in DRAM.

The hard disk's competence is bulk storage for the PC.

The personal computer consumes 30 percent of the worldwide production of semiconductor components. The PC helped the semiconductor market grow from $14 billion in 1981 to $188 billion in 2002. The PC is performance oriented, but, at introduction, the PC's performance was woefully inadequate. Its makers strove to improve it. As long as its performance wasn't good enough, buyers paid premiums for leading-edge systems. It took twenty years, but the PC's performance now exceeds the needs of most buyers. Leading-edge PCs still offer leading-edge performance at premium prices. But value PCs now offer good-enough performance at competitive prices.

The value PC is shifting engineering design emphasis from tethered systems to untethered systems.

## Supply and demand

The PC story illustrates supply and demand. At its debut, PC customers were early adopters and PC performance was well below their desires. Over time, the PC's performance improved at a rate close to the Moore's-law rate of improvement of its components. The expectations of many customers rose, but the PC's customer base also broadened to include less-demanding customers.

There's no necessary connection between the rate of growth in PC performance (supply) and the rate of growth in the demand for performance. PC performance grew faster than customers' demands.

The hard disk also illustrates the independence of rates of improvement in supply and in demand. The supply of hard-disk capacity improved rapidly (60 percent per year); the demand for capacity grew more slowly. When the capacity of 5¼-inch hard disks overshot demand, the market switched to lower-capacity, cheaper 3½-inch hard disks. As the capacity of 3½-inch hard disks overshoots demand, the market will shift to 2½-inch hard disks.

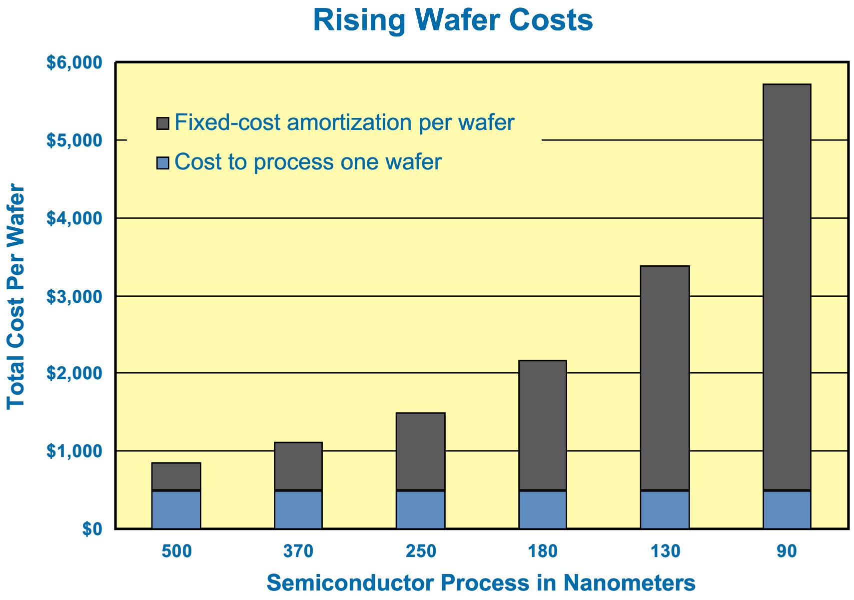

**Fig. 1. Amortization of fixed costs for buildings and equipment now dominates total cost to process a wafer.**

## Semiconductor fabrication

**Scaling in *Flatland*: the theory.** Here's how Moore's law works. Transistors are two-dimensional structures. Shrinking each dimension of a transistor by a factor of 1.41 (the square root of two) every eighteen months fits twice as many transistors on a wafer. That translates to the same number of chips, each with twice as many transistors, or to twice the number of chips, each with the original number of transistors. The cost is the same because the cost to process a wafer depends only on its size. That's the theory. And that's been the way it's worked for a long time.

A semiconductor plant has two kinds of costs: fixed and variable. The fixed cost is the cost of the buildings and equipment. The variable cost is the cost to process a wafer. The total cost of a wafer is its variable cost, plus its share of the fixed cost. Plus, each chip has to pay its share of the mask cost. Masks hold the patterns for a *specific* chip design. A million-dollar mask set adds one dollar a chip to a million-unit run and adds one thousand dollars a chip to a thousand-unit run.

**Scaling: the practice.** If the plant costs (fixed costs) and the mask costs (design-specific costs) are small relative to the wafer-processing cost (variable cost), then Moore's law in practice matches Moore's law in theory. Unfortunately, the fixed costs for buildings and equipment have been doubling with each process generation. The buildings and equipment for 130-nm transistors cost about $2 billion, while the buildings and equipment for the next-smaller 90-nm transistors cost about *$4 billion.* Mask costs are rising even faster. Masks for 130-nm transistors cost about $650,000, *per chip design.* Masks for 90-nm transistors may cost $1.4 million.

Fig. 1 shows how the amortization of fixed costs has become more significant than the cost to process a wafer (variable or operating costs). Fig. 1 makes the situation look more serious than it is because fig. 1 doesn't show the value of shrinking the transistors. With each process generation, it takes half as many wafers to produce the original number of transistors. Fig. 2 accounts for this by normalizing the cost to the cost of a transistor in a 500-nm process. In fig. 2, old processes (500 nm, 370 nm, and 250 nm) are fully amortized. These figures do not consider mask costs. Escalating mask costs make smaller transistors even more expensive. Fig. 2 shows that a 250-nm process makes the cheapest transistors.

The fabrication plant's fixed costs are amortized over some period, such as three years. Once the plant is paid for, it produces cheaper chips than a fab burdened by amortization costs.

All of this; rising fixed costs, rising mask costs, and fully amortized fabs; leads to an interesting result. *The cheapest chips may not be the ones with the smallest transistors.* There's a "value transistor" just as there's a value PC.

For chips up to 25-million transistors (about half the complexity of a Pentium 4 microprocessor), the cheapest transistors come from a 250-nm process—three manufacturing generations behind the leading edge. That's true for production runs to at least eight million chips. For longer runs, smaller transistors may be cheaper because the higher fixed costs can be spread over more chips. For short runs, the mask cost may determine the application's cheapest process. The newer semiconductor processes at 180 nm, 130 nm, and 90 nm produce smaller, faster transistors, but the transistors cost *more* than ones from older processes.

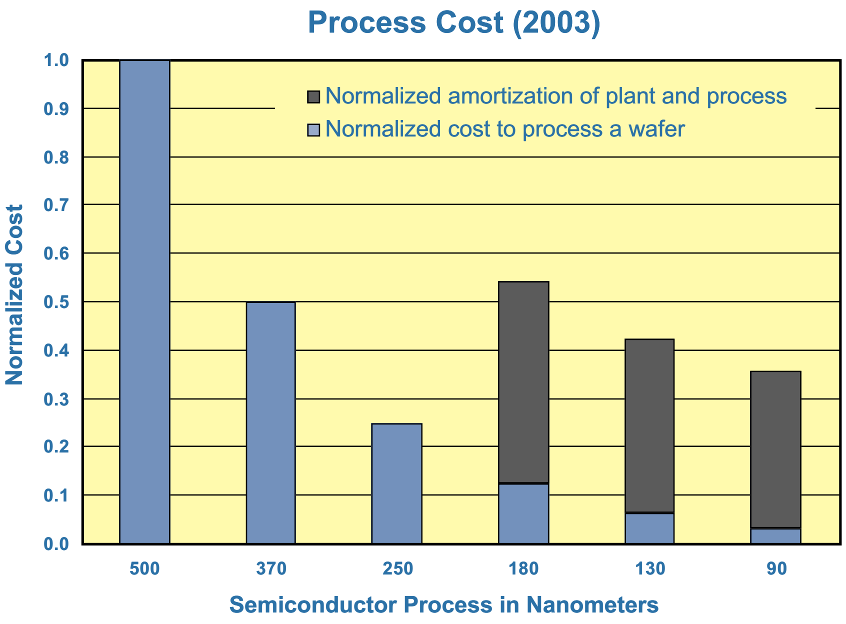

**Fig. 2. Fully amortized 250-nm fabs make the cheapest transistors.**

**The value transistor.** In the early days of integrated circuits, transistors were large. Like the early days of the PC, when the PC's performance didn't satisfy demand, the integrated circuit's big transistors didn't meet all the needs of *any* application. But, transistors got smaller and faster. In time, a demand profile appeared: not everyone chose the smallest and fastest. *Now, for a wide range of applications, the available transistors are good enough.* The value transistor is the cheapest transistor that's good enough. A smaller, faster transistor might work, but designers won't pay extra for it.

As more applications are satisfied with available transistors, fewer applications are left to pay for new buildings and equipment, for new manufacturing processes, and for the more-expensive, design-specific masks that go with smaller transistors. (Replay the value PC story with "smaller transistors" substituted for performance.)

The Moore's-law rate of progress has been business as usual in semiconductors. That was fine when the transistor wasn't good enough. Now, moving to the next-smaller transistor becomes more expensive with each generation even as fewer applications are willing to share the costs. Because making improvements the way the industry knows how to make them is safer than trying something new, the semiconductor industry has continued to march toward smaller transistors. This inertia has carried the industry past the point of payoff for new ideas such as 3D-integrated-circuit fabrication.

The value transistor, like the value PC, evolves. Next year's value PC will have better features and more performance than this year's value PC. Next year's value transistor will be smaller and faster than this year's. When the 180-nm fabs are paid for, the value transistor will snare applications for chips with fifty million transistors.

Fig. 2 makes it look as if the value transistor is produced by the fab with the most recent process for which the fixed cost has been paid. But chip volume and mask cost determine the answer in a particular case.

**Falling adoption rates.** This "value transistor" theory sounds interesting. Is there any evidence to support it? There is, if we look at semiconductor process (e.g., 250 nm, 180 nm) adoption at foundries.

Two kinds of companies build chips: integrated device manufacturers and foundries. Integrated device manufacturers own the bulk of chip production today, but the foundry business is growing as a percent of the total. Integrated device manufacturers are what they sound like; they are vertically integrated from chip design through manufacturing and sales. Integrated device manufacturers include **Intel** (INTC), **Motorola** (MOT), **STMicroelectronics** (STM), and **Texas Instruments** (TXN). In the integrated device manufacturers, process transitions are by corporate fiat. Process transitions are built into their business models, so they don't measure demand. By contrast, foundries are demand driven. Their customers are the fabless semiconductor companies such as **Altera** (ALTR), **ATI Technologies** (ATYT), **Nvidia** (NVDA), and **Xilinx** (XLNX). **IBM** (IBM) is both an integrated device manufacturer and a fab.

Foundries offer the semiconductor processes their customers demand. The adoption rate is the increase or decrease in the percent of wafer starts for a particular semiconductor process. If we plot adoption rates by process, two trends are evident. First, the adoption rates for newer processes are falling. Five or six years ago, when a foundry introduced a new process, customers adopted it quickly. By the end of the first year, a new process might account for *30 percent* of wafer starts. Now, recently introduced processes capture but a *few percent* of wafer starts by the end of their first year. Second, *old processes now hold their own for years*, as a percentage of wafer starts, even as the total number of wafer starts grows rapidly. Both trends are long-term trends that are independent of the semiconductor industry's boom and bust cycles.

The longevity of old processes shows that many applications have found their value transistor—it isn't cost effective to move to a new process. Falling adoption rates for new processes show that fewer applications are willing to pay a premium for smaller transistors.

## Small consequences

A funny thing happened on the way to 22 nm. Following the *International Technology Roadmap for Semiconductors* (http://public.itrs.net) to its forecast horizon takes us to 22-nm transistors. We won't be going there; that forecasts the no-sea-change scenario.

With the design objective shifting to cost-performance-per-watt, smaller transistors may not just be more expensive, they may not be better at all. As the transistor gets smaller, it gets faster and it uses less power *to do work*, but it leaks more power when not doing work. Leaky transistors are bad for devices that care about battery life.

If the transistors in your system are working all the time, you want the smallest transistors you can get because they have the lowest active power. If the transistors in your system are mostly idle, you want bigger transistors because they leak less. As more transistors fit on a chip, fewer of them are active and more are idle (lots of functions with few in use at a time). Most applications fit between the extremes; there's a transistor that's right for them. And more and more, it's not the transistors made with the leading-edge process.

## New memory and storage

Flash memory, SRAM, DRAM, and the hard disk are unsuitable for untethered devices because they were developed in watt-rich environments. Memories' Holy Grail has the non-volatility of flash memory, the speed of SRAM, and the density of DRAM. Up to now, no candidates could compete against incumbents because the incumbents occupied secure economic niches in the PC. Emerging untethered applications provide new incentive for R&D investment and new technical niches that the PC's components cannot satisfy.

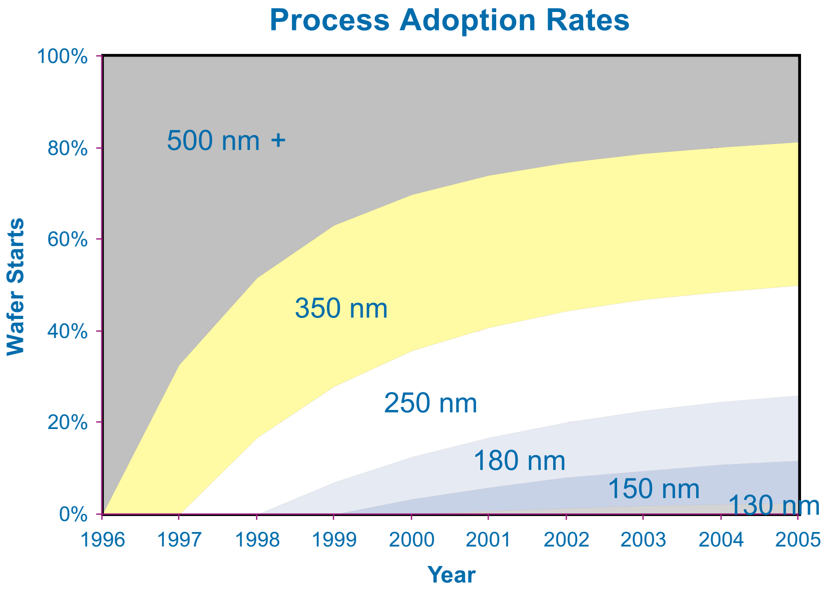

**Fig. 3. As the fab moves to new semiconductor processes, *adoption rates decline.***

**Non-volatile memory**. The leading non-volatile memory candidates for untethered applications are ferroelectric random-access memory (FRAM), magnetoelectric random-access memory (MRAM), and ovonic unified memory (OUM) (*Dynamic Silicon*, May 2002). Each of these has impressive backers and each has been slow to develop because there was no hope of displacing flash memory's economic niche in the PC. The newcomers couldn't compete on cost with any of the high-volume, entrenched incumbents.

But non-volatile memory contenders now have applications, in untethered devices, where they are the best solution. These promising candidates include latecomers such as Axon Technology's programmable metallization cell memory (PMCm) (*Gilder Technology Report*, March 2003).

**Storage**. There's a similar story for the hard disk. An untethered device needs storage, but a hard disk's size, startup delay, and power consumption render it unsuitable. This is an opportunity for MEMS-based storage (*Dynamic Silicon*, May 2001); it couldn't compete with the hard disk's entrenched position in the PC.

Programmable logic devices (*Dynamic Silicon*, Special Report "Dynamic Logic vs. Computing") are generic chips that are customized in the field. Programmable logic devices really need non-volatile memory. Non-volatile memory will increase programmable logic devices' speed, security, circuit capacity, and energy efficiency (*Gilder Technology Report*, March 2003), making them a practical substitute for microprocessors and for digital signal processors.

## 3D-integrated-circuit fabrication

3D-integrated-circuit fabrication is a new, backwards-compatible approach to chip making.

3D-integrated-circuit fabrication *stacks wafers* to grow chips *vertically*. Hundreds of thousands of vertical connections per chip carry signals "upward" on wires that are much shorter than the distances across a chip. Shorter wires mean faster signals at lower power. Amplifiers to drive the shorter wires can be smaller, which saves additional power and area.

But 3D integrated circuits are not just a density play; stacking wafers has other advantages. Stacking enables *mixing analog-circuit wafers, digital-logic wafers, and memory wafers*. Each wafer type can be built in its own semiconductor process before stacking. This reduces wafer-processing cost and complexity.

Because of the on/off nature of binary logic, digital circuits scale well to smaller transistors. Analog logic's transistors operate in the "linear" region between off and on. Because analog circuit operation is *very* sensitive to transistor size, scaling is difficult and labor intensive. Analog circuits sit between the digital logic and the physical world on the input and output paths. These circuits don't need to scale as rapidly as digital circuits because their physical interfaces don't change. If analog circuits are stuck on the same chip with digital logic, then as the digital logic scales to faster circuits there's enormous overhead in scaling the analog circuits at the same time. Building the analog circuits on their own wafers and the digital circuits on their own wafers permits the slow evolution of analog circuits and the fast evolution of digital circuits.

The story is similar for mixing digital logic and memory. Memory cells want low-leakage transistors, while logic circuits tolerate leakage to get speed. This is what complicates building chips that mix memory cells and digital logic. Stacked wafers can mix process characteristics (e.g., high-speed transistors on one wafer and low-leakage transistors on another) and even *process generations* (e.g., 250 nm, 180 nm, 130 nm). Individual wafers in a stack are thinned so that the resulting chips are no thicker than conventional "2D" chips.

Chip makers have been stacking *individual chips* for a long time for special applications. Compared to 3D-integrated-circuit fabrication, these efforts are clumsy and expensive.

For foundry *customers*, 3D-integrated-circuit fabrication should be a cost-effective alternative to moving to the smaller transistors of the next semiconductor-process generation. For the *foundries*, 3D-integrated-circuit fabrication is *far cheaper* than the billions of dollars necessary to produce the next-smaller transistor. It increases capacity and performance over 2D chips built of the same transistors. 3D memory chips, for example, might have four times the capacity and twice the performance of 2D chips built in the same process.

There's no significant technical barrier to 3D-integrated-circuit fabrication today. Tachyon Semiconductor and others have demonstrated the feasibility of 3D-integrated-circuit fabrication. The barriers are in habit (forty years in *Flatland*) and in licensing. The first company to demonstrate its usefulness in commercial products will open the floodgates.

## The value fab

The value transistor decreases the incentive for advanced process development. If transistors are good enough for most applications, then the fabs are good enough too. The value transistor and the demands of emerging economies for non-demanding products lead to the "value fab." Value fabs build products for emerging economies and they build products with value transistors. These fabs buy value processing equipment. Value fabs will reduce margins in semiconductor processing equipment in the same way that the value PC reduced margins in the PC business.

Building products with value transistors means there's little to be gained in shrinking transistors and in developing processes. These fabs will explore lower-cost alternatives, such as 3D-integrated-circuit fabrication.

The proliferation of value fabs will create demand for value processing equipment. The cost of fabs, which has been escalating exponentially with shrinking transistors, will level off and decline. It'll be a difficult time for equipment suppliers, such as **Applied Materials** (AMAT) and **KLA Tencor** (KLAC), that have been supplying leading-edge equipment at premium prices.

When will we see value fabs? They're here now. The foundries *are* value fabs. Foundries are demand driven; they build the transistors their customers want. Fig. 4, which is adapted from a TSMC chart, shows relatively steady demand (in wafer starts) for 350-nm and for 250-nm processes. The customers buying chips in these old processes buy them because the financial and technical incentives aren't sufficient to entice them to smaller transistors. Customers with performance-oriented requirements pay a premium for leading-edge transistors, but fewer and fewer pay these premiums over time.

## The big picture

The semiconductor industry has been operating on a shrink-the-transistor formula for forty years. Shrinking the transistor has been successful because the goal was cost performance and because the transistor wasn't good enough. The coming change in semiconductors reflects the change in direction from cost performance to cost-performance-per-watt and transistors becoming good enough for most applications. Cost-performance-per-watt systems must balance active power and leakage power, so smaller, leakier transistors aren't always better. This calls for new system-component design.

The semiconductor industry is replaying the development of the automobile: mass production from Model Ts (every driver was a mechanic) to muscle cars (cars performed more than well enough), followed by a difficult transition to energy-efficient cars without sacrificing performance. Cars had to be reengineered, not just restyled over bigger engines.

## The next phase

With the PC as king of electronics platforms, the members of its court ruled markets. As the PC's importance declines, the crown is passing to untethered applications. The cell phone, not PDAs, will be the jumping-off point in the market struggle to define the new architecture and the pieces comprising it. These new system components will form out of the transistor wake already left by Moore's-law progress.

The PC reaching its value plateau and the rise of untethered devices provide incentive for new R & D investment.

Microprocessors and digital signal processors will be de-emphasized in favor of more efficient, more direct implementations that will use a new generation of non-volatile programmable logic devices.

Non-volatile memories will leverage their fifteen-year history of development. They were unable to gain a foothold against the PC's entrenched memory components. New non-volatile memories, coming from the established ranks of the FRAM, MRAM, and OUM developers or from startups such as Axon Technology, will enable the proliferation of untethered applications.

Non-volatile memory will replace SRAM in today's programmable logic devices. This will make programmable logic practical—energy efficient—in untethered devices.

Emerging untethered applications will demand great numbers of miniature, integrated sensors and actuators, which will lead to the proliferation of microelectromechanical systems. The unsuitability of hard disks for untethered applications creates investment incentive for development of MEMS-based non-volatile storage.

For cost-performance-per-watt devices, Moore's law no longer gives us the best transistors, only smaller ones—smaller transistors that cost more and that are energy inefficient. Three-D-integrated-circuit fabrication will increase density without having to use smaller, less-efficient transistors. Three-D-integrated-circuit fabrication also enables mixing process generations and circuit varieties (analog, memory, and logic).

(For over a decade, the *least-demanding* PC buyer was dissatisfied with memory capacity and with microprocessor performance.) The microprocessor's drive for speed led the industry to equate Moore's-law progress and semiconductor industry progress. The industry's planners and analysts thought that if Moore's-law progress stopped, the semiconductor industry would stagnate. It's not true. We shall see, as the industry shifts from cost performance to cost-performance-per-watt, that there's plenty of room to progress. It's just not where the progress has been for the last twenty years with the PC, or for the last forty years with the shrinking transistor.

The emergence of the value transistor—a range of transistor sizes that suits most applications—means that what engineers use transistors for becomes more important than how small they are. The value transistor will dominate transistor demand, and it will change the semiconductor-processing equipment companies from premium-priced, leading-edge equipment suppliers to value equipment suppliers. The combination of transistor use being more important than transistor size and the dominance of the value transistor eliminates the competitive advantage of the integrated device manufacturers over the foundries, accelerating the fragmentation of integrated device manufacturers. (The ability to make smaller transistors will no longer be the competitive advantage.)

This isn't gloom and doom; this is *a sea change*. I see vast opportunity in untethered devices. I see great possibilities in being free of shackles that slaved the industry to shrinking transistors. It's time to unleash our creativity, to use transistors rather than to just harvest the byproducts of making them smaller. The processing requirements of future untethered devices dwarf the capabilities of today's most powerful PCs. The future's untethered devices will do all the things today's PCs do, but they will know where they are, how they are moving, what their environment looks like, and who's using them. They'll need continuous, agile broadband connectivity, vastly improved human and physical-world interfaces, and enormous computational power. They'll have awareness and responsibility that's unimaginable in today's PC.

## System Components

| | FOR TETHERED SYSTEMS | | | FOR UNTETHERED SYSTEMS |
|---|---|---|---|---|
| DESIGN OBJECTIVE | COST | PERFORMANCE | COST PERFORMANCE | COST-PERFORMANCE-PER-WATT |
| Processing | Microprocessors, microcontrollers | Microprocessors, digital signal processors | Microprocessors, digital signal processors | New |
| Memory | Flash, DRAM | Flash, SRAM, DRAM | Flash, SRAM, DRAM | New |
| Storage | — | Hard disks | Hard disks | New |
| Applications | Embedded systems | Workstations | PCs | Cell phones, personal digital assistants, digital cameras, MP3 players, GPS receivers |
| Comments | Continued growth | Overtaken by PCs | Diminishing returns, value PCs | New investment, emerging, rapid growth |

Fig. 4. Current and future components of electronic systems.

## Sea Change Scorecard: Who Wins, Who Loses

| COMPANY | TYPE OF COMPANY | FUTURE POSITION | THE WAY I SEE IT |
|---|---|---|---|
| Altera, Xilinx | Fabless | Excellent | The transition in design objective from cost performance to cost-performance-per-watt is an opportunity for new products from established programmable logic companies. Non-volatile memory for programmable logic devices will greatly increase demand. |
| ARM | Fabless | Excellent | ARM has established an early lead and a strong position in soft-core micro processors for untethered applications. |
| GSMC, SMIC | Foundry | Excellent | GSMC, SMIC, and other Chinese foundries benefit from rapid growth in domestic demand and also from the shift from integrated device manufacturers to foundries as more applications reach their value transistor. |
| Legend | Computer Systems | Excellent | While the shift from leading-edge PCs to value PCs will squeeze most PC makers, Legend is at the center of a fast-growing market in value PCs. |
| ARC International, Tensilica | Fabless | Good | ARC and Tensilica offer microprocessors that are customized to be efficient for a particular application. The "customized microprocessor" concept sells better than a direct transition from a fixed-instruction-set microprocessor to programmable logic. |
| Ascenium, GateChange, QuickSilver | Fabless | Good | Rapid growth in untethered applications should mean rapid growth for companies that build efficient implementations. The opportunity for these companies should be excellent except that they have to overcome entrenched opposition, and they have to reeducate potential customers. |
| Chartered, TSMC, UMC | Foundry | Good | As more applications reach their value transistor, more business shifts from integrated device manufacturers to foundries. The opportunity for TSMC and for UMC would be excellent except for investment restrictions by the government of Taiwan. |
| Elm Technology, Tachyon Semiconductor | Fabless | Good | Companies that license and develop 3D-integrated-circuit fabrication should do well as their techniques and products are adopted. |
| MemoryLogix, Transmeta, VIA Technologies | Fabless | Good | The microprocessor may not be the workhorse in future systems, but it will still be there in a supervisory role. As software content rises in electronics, the x86 will increasingly invade embedded systems. MemoryLogix, Transmeta, and VIA have the opportunity to offer x86 as a low-end chip or as a soft core. |
| Dell Computer, Gateway, Hewlett Packard | Computer Systems | OK | The shift in emphasis from leading-edge PCs to value PCs will squeeze margins for PC makers. |
| Applied Materials, KLA Tencor, Lam Research, Novellus | Semiconductor Equipment | Struggle | As more applications reach their value transistor, fewer chip makers need leading-edge equipment. The semiconductor processing equipment industry will shift from profitable, high-end processing equipment to low-margin, value equipment. |
| Intel, Motorola, Texas Instruments | Integrated Devices | Struggle | The appearance of the value transistor decreases the competitive advantage of owning an integrated fab. |

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.

## Got Questions?

**Visit our subscriber-only discussion forum, the Telecosm Lounge,
with George Gilder and Nick Tredennick, on www.gildertech.com**

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<br>Brion Shimamoto | Bret Swanson | Mary Gorski<br>John Hammill | Marie Lavinio | Charles Bork | Rosaline Fernandes |

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