Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 12 | The Investor's Guide to Breakthrough Micro Devices | December 2002

Moore's Law and the Real World

Invited commentary from the CEO of a semiconductor startup: "For 30 years semiconductor growth has been driven by the exponential improvements in the cost/performance ratio of semiconductor devices, a phenomenon known to us all as Moore's law" (*EE Times*, 2 December 2002, pg. 45).

This quote echoes widespread belief that Moore's law is the engine of growth for the semiconductor industry, and there's concern that semiconductor processing may be reaching limits that block Moore's-law progress. Since people believe that Moore's law drives the industry, they believe that as Moore's law goes, so goes the industry.

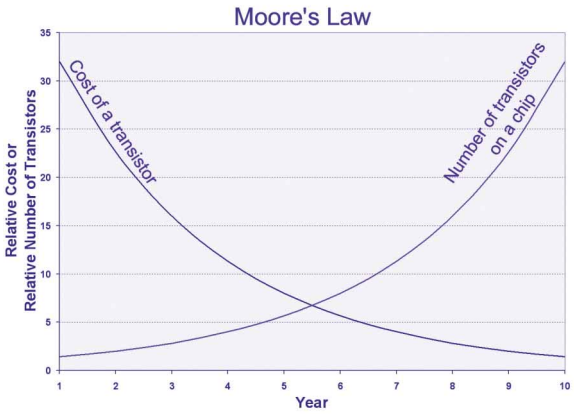

Fig. 1. Moore's (supply) law: transistors shrink so chips get smaller or they hold more transistors.

Fig. 1 shows the complementary effects of Moore's law—what would be better termed Moore's *supply* law—in semiconductor manufacturing. Moore's-law progress shrinks transistors. As the transistors get smaller, existing designs fit on smaller chips. More chips fit on a wafer, which is the unit of processing. (Wafers are the measure of manufacturing capacity.) The cost to process a wafer doesn't vary with the size of its chips. Smaller chips are, therefore, cheaper. Cheaper chips enable more cost-sensitive applications. That's one effect of Moore's law: the cost of a fixed number of transistors drops by half in eighteen months. As transistors get smaller, more fit on the same chip. That's the second effect: twice as many transistors occupy the same space in eighteen months. More transistors enable uses out of the reach of the previous-generation chip.

Moore's-law progress enables more cost-sensitive applications at the low end, and it enables greater functions and higher performance at the high end. *Enable* here means to make *possible* or to make *feasible, if demand is there*.

I'll take you through the theory of Moore's law, starting with basic assumptions, and I'll inject reality in steps to bring you to where we are today and say what it means to the industry.

Most of today's semiconductor production is with 200-mm-diameter wafers, so I start with that assumption. I'll assume we're building a 25-million-transistor chip (about half the complexity of a leading-edge Pentium). How much will this chip cost? What will Moore's law do for the chip's cost as the semiconductor process advances?

More advanced semiconductor processes make smaller transistors. I'll build imaginary plants for six transistor sizes: 500 nm, 350 nm, 250 nm, 180 nm, 130 nm, and 90 nm. Today's leading-edge process is 130 nm, with leading suppliers beginning their move to 90 nm. The cost of making chips divides into fixed costs and variable costs. Fixed costs build the plant and furnish it with equipment. There are also fixed costs for developing the semiconductor process, for the software tools to build the design, and for the mask set that repre-

In This Issue:

Wood, paper, metals, cloth, and plastics are good enough for a wide range of uses. Now, even silicon joins the list. From now on, what we do with transistors is more important than how we make them. But the semiconductor industry has thirty-five years of momentum developing and refining the silicon medium, and it is not about to give up making transistors smaller, whether it makes economic sense or not. The latest semiconductor processes cannot count on amortizing their costs across all market segments—a new phenomenon. I could have called this issue *Moore's Law Meets Market Segments*. Who would have thought that Moore's law might end, not from technical difficulty, but from declining market interest?

sents the chip. (The difference between the semiconductor equipment and the process is like the difference between the equipment and the detailed recipes-and-procedures in an industrial kitchen. Masks are like photographic negatives, defining a chip's details in a layer-by-layer buildup.) The plant processes chips in "boats"—plastic tubs of twenty-five wafers. Running a wafer through the plant is a variable cost, about $500 a wafer.

I assume each plant processes 25,000 wafers a month (this is a median figure) and that it amortizes fixed costs

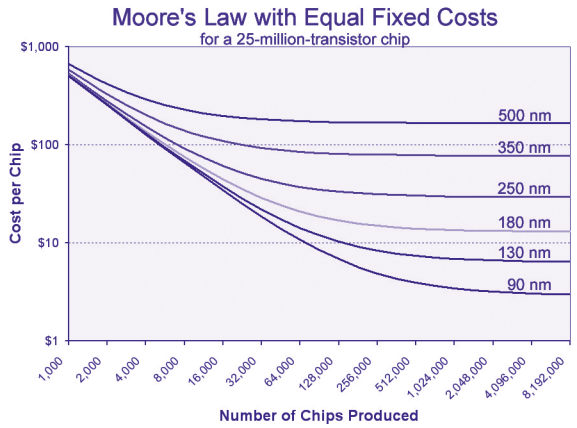

Fig. 2. If fixed costs were the same, shrinking transistors from 350 nm to 130 nm would drop the cost of a 25-million-transistor chip from $78 to $7.

for the plant, for the equipment, and for process development over the first three years of production. If the plant doesn't run at capacity (a frequent occurrence lately), the amortization period will likely exceed three years. Amortizing costs over four years doesn't change the qualitative values, but makes differences more difficult to see in small figures. The fixed cost for a mask set is amortized over the number of chips produced using that mask set. Fig. 2 shows how the cost per chip varies by semiconductor process and by the number of chips built for a 25-million-transistor chip.

The cost per chip is high for building a few thousand chips because the fixed costs (plant, equipment, process development, and mask set) dominate. At low volumes, there's little cost advantage in making smaller transistors. Variable costs dominate for production runs of millions of chips.

Plants, masks, and equipment

Fig. 2 is Moore's law in theory. It assumes the *same* fixed costs across all semiconductor processes. That's not the situation in the real world. Smaller transistors are harder to make. The plant and equipment to build finer geometries are more expensive. Developing processes for smaller transistors is more expensive. Mask sets for smaller transistors are more expensive. The cost of the plant and equipment approximately doubles with each new process generation. Process development escalates more slowly and mask cost escalates a little faster, but all the fixed costs rise as transistors get smaller. Fig. 3 factors in escalating costs for the plant, for the equipment, for process development, and for the mask set.

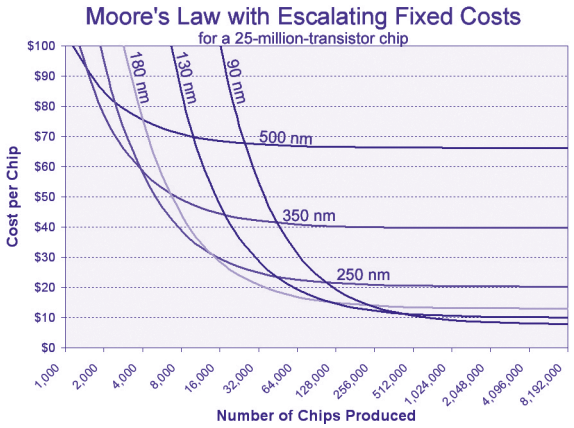

Fig. 3. Smaller transistors are cheaper than large transistors, but the advantages of shrinking shrinks with each generation.

In fig. 2, the curves don't cross; in fig. 3, they do. If fixed costs are all the same (fig. 2), it's always cheaper to build smaller transistors. If fixed costs rise as transistors get smaller (fig. 3), then for production runs of a few thousand chips (where fixed costs dominate chip cost) it's cheaper to build larger transistors. Now that we have accounted for escalating fixed costs, two important characteristics of these costs appear in fig. 3.

First, the cost-crossover point between generations is *moving to the right*. For a 25-million-transistor chip, 350-nm transistors become cheaper than 500-nm tran-

DynamicSilicon

Editors Nick Tredennick, Brion Shimamoto
Art Director Charles Bork
Managing Editor Marie Lavinio
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.

sistors at 1,500 chips, but 90-nm transistors don't become cheaper than 130-nm transistors until production runs exceed 500,000. For larger chips, crossover points move toward smaller production runs; for smaller chips, crossover points move toward larger production runs. Crossover between 90-nm transistors and 130-nm transistors for a 2-million-transistor chip requires production runs of millions of chips. Advanced processes want big chips *and* large production runs. Moving to larger wafers, from 200 mm to 300 mm, pushes crossover points toward larger production runs.

Second, the chip's *cost advantage decreases as transistors get smaller*. For large production runs, moving from 350 nm to 250 nm reduces chip cost from $40 to $20—a saving of $20 or 50%. For a production run of eight million chips, moving from 130 nm to 90 nm reduces chip cost from $10 to $8—a saving of $2 or 20%.

Escalating fixed costs push the economic crossover point to the right. If you are designing a 25-million-transistor custom chip as the brains of an espresso machine, the chip will be cheaper in a 130-nm process than in a 90-nm process unless you expect to ship more than 500,000 machines. If the custom brain is fewer than 25 million transistors, you will need to sell even more machines.

Escalating fixed costs also mean that, for a given number of transistors, the incentive to move to smaller transistors diminishes with each generation. The incentive diminishes both in absolute dollars per chip and as a percent of the cost of the chip. If you are already shipping the espresso machine with a custom chip designed in a 180-nm process, each chip is costing you $13. Your engineers could cost-reduce the design by moving the brain-chip to a 130-nm or to a 90-nm process. You would save about $2 per chip. If the bill of materials for the espresso machine totals $1,000, a $2 savings isn't much incentive. It would also cost a few engineers and several months. You would have to sell another 500,000 machines just to break even on the move to a new process.

Old processes, new processes

Fig. 4 shows what happens to chip cost with the depreciation of fixed assets. In fig. 4, I assume the three leading-edge processes (180 nm, 130 nm, and 90 nm) are being built in plants that are not fully depreciated. The cost of leading-edge chips, therefore, includes the amortization of the plant and its equipment across 25,000 wafers a month over three years of operation. The three trailing-edge processes (500 nm, 350 nm, and 250 nm) are being built in plants that are fully depreciated. The cost of trailing-edge chips, therefore, does not include dollars for plant and equipment amortization.

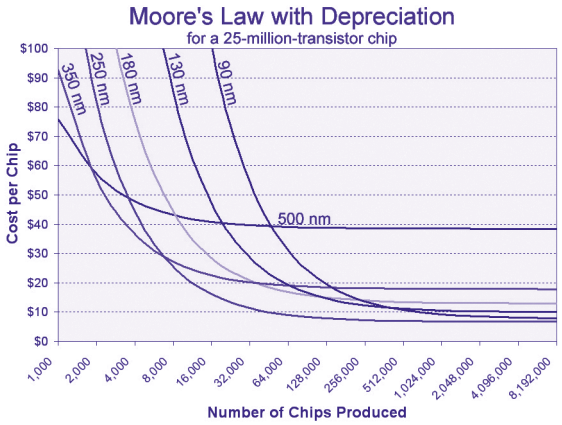

Fig. 4. Here, foundries for 180-nm, 130-nm, and 90-nm chips still amortize fixed costs into the cost of a 25-million-transistor chip. Fully depreciated foundries for 500-nm, 350-nm, and 250-nm offer lower-cost chips.

I've added another dose of reality; this time the result is even more surprising. The cost to build a 25-million-transistor chip in a newer process (180 nm, 130 nm, or 90 nm) never drops below the cost to build the same chip in a 250-nm process! A fully depreciated 250-nm process turns out the cheapest 25-million-transistor chips. As long as cost is more important than performance, the big transistors on the old process beat the small transistors on newer processes. Once the 180-nm foundry is fully depreciated, it will turn out the cheapest 50-million-transistor chips (Moore's law).

Applications span a performance gamut from hair dryers, washing machines, and blenders to computers, video games, cell phones, and set-top boxes. The electronics market consumes billions of microprocessors each year. Most of these are four- and eight-bit microcontrollers. These applications are cost-oriented and are not performance-oriented. One or two million transistors make a very capable microcontroller. Twenty-five or fifty million transistors are more than enough for a *huge* range of applications.

Leading-edge applications pay a premium for performance and foot the bill for new processes. But with each new process generation, the range of applications available to pay premium prices shrinks.

Small chips

Fig. 4 showed cost curves for leading-edge processes and for trailing-edge processes for a 25-million-transistor chip. Fig. 5 shows the curves for a 2-million-transistor chip. The smaller chip changes the scale; small chips are less than a tenth of the cost of large chips, but the conclusions remain the same. The fully depreciated 250-nm process builds the cheapest small chip. The 2-million-transistor chip's equal-cost crossover points have moved to the right by about ten times the number of chips at the 25-million-transistor

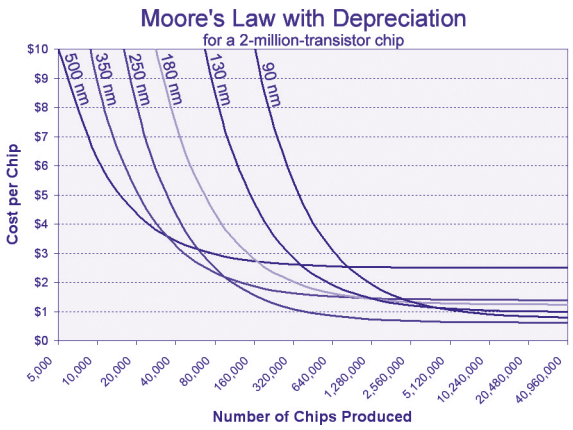

Fig. 5. Foundries for 180-nm, 130-nm, and 90-nm chips amortize fixed costs into the cost of a 2-million-transistor chip. Fully depreciated foundries for 500-nm, 350-nm, and 250-nm offer lower-cost chips.

chip's crossover point. For all reasonable production volumes above 100,000, the 250-nm process makes the cheapest chips. For production volumes below 100,000 units, it would be cheaper to produce the 2-million-transistor chip on an even older process with bigger transistors.

Connection-limited chips

For communication with the outside world, chips can have a ring of wire-bonding pads, or they can have an array of solder balls. Whether it's a pad ring or a (more area-efficient) solder-ball array, these connections are a constraint on chip area. The area of the circuit's transistors is one limit on making the chip smaller, and the area to accommodate the chip's external connections is a second limit. If the connections limit chip size, making the transistors smaller doesn't make the chip cheaper. Fig. 6 shows the cost of a 25-million-transistor chip that becomes connection-limited at 350 nm. Once the chip's size is limited by the area for its connections, there's no cost savings to move to a more-advanced process.

Each process generation doubles the number of transistors on a connection-limited chip, so each generation takes another bite out of the application space. The bite that's taken out is for cost-oriented chips that will never need smaller transistors.

Supply and demand

I feel like a broken record, bringing up supply and

The Value Transistor

What's a "value" transistor? A value transistor is the right transistor for the job.

The value transistor is a recent phenomenon. We're used to thinking that Moore's law makes the transistor better with each process generation. This is no longer true for every application. Each process generation is now creating what I call the value transistor—one that is good enough to completely satisfy a set of applications. This is analogous to what happened with the personal computer—"value PCs" now satisfy a growing segment of users. When I started thinking about how to describe the value transistor, I thought I would hit upon the perfect description to convince you that there is such a thing. I built a mathematical model and I tried ways to demonstrate its existence. It turns out that it's complicated, so there's no one description that illustrates the value transistor for all cases. One way to identify a value transistor looks at power use in chips for mobile applications.

The most common transistor—the CMOS field-effect transistor—dissipates energy in two ways: as *active power* and as *leakage power*.

Active power switches the transistor on or off. Active power depends on the size of the transistor and on its frequency of operation. Transistors burn active power while doing arithmetic or while making control decisions. Big transistors use more power (the bigger area of a large transistor requires more electrical charges). Switching the transistor faster pumps more charges from the power supply to the transistor and then to ground.

Leakage power is the power dissipated as electrical current leaks when the transistor is not being switched. Big transistors don't leak much. Small transistors, with their thin conductors and thin insulation, leak more.

As transistors get smaller, the active power decreases and the leakage power increases. This means that for some applications, there is a *value transistor*: make it larger and its active power is too high; make it smaller and it leaks too much. Fig. 8 shows curves for a 200-million-transistor chip operating at frequencies from 50 MHz to 400 MHz. Note how specific the fig. 8 example is. It is for a particular number of transistors operating at a particular range of frequencies.

The percentage of transistors that are active (being switched on or off), as opposed to idle (not being switched), heavily influences the shape of the curves.

Active power isn't straightforward to calculate. It depends on the operating frequency and on the percent of transistors that are active. In a 1-Mb memory chip, for example, only a tiny percentage of the transistors are active, while the rest sit idle, holding information. When only a few thousand transistors fit on a chip, designers got all they could out of every transistor, perhaps 10% of the transistors were in use at any given time, while 90% sat idle. Today, with tens or hundreds of millions of transistors on a

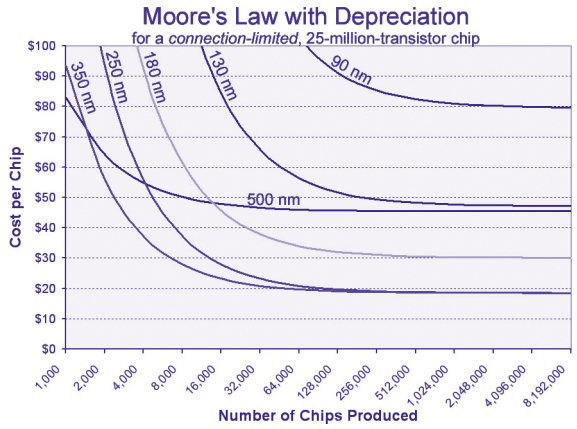

Fig. 6. The 25-million-transistor chip that is connection-limited in a 350-nm process won't get cheaper in a more-advanced process.

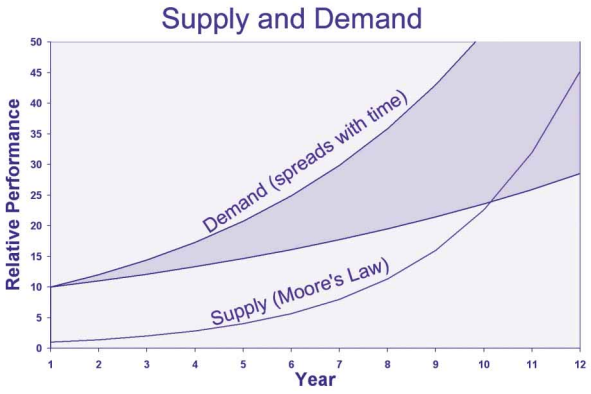

Fig. 7. The supply of transistors on a chip grows with Moore's law. Demand for transistors is difficult to measure and it spreads with time.

demand so often. I'm doing it because there's not enough recognition of how important it is. Fig. 7 shows supply and demand curves for the PC. This explanation of supply and demand illustrates how the market began its shift from buying performance-oriented PCs to buying value-oriented PCs. In this case it's crucial to understanding the semiconductor business, because I postulate the "value transistor."

When the PC was introduced in 1981, its performance didn't satisfy anyone. Demand for performance was ten or a hundred times what the PC delivered. Intel and other companies supplying the PC's microproces-

chip, it wouldn't be unusual for only 1% of the transistors to be active at any given time, while the remaining 99% sit idle. For "always-on" devices, only a tiny percentage of the transistors might be active while the device waits for input. Active power is the product of the transistor's switching energy times its frequency of operation times the number of active transistors on the chip.

All of the transistors on the chip leak all of the time, so leakage current is straightforward to calculate. For example, at 130 nm each transistor leaks about 5.6 nanowatts. That doesn't sound like much, but that means a 200-million-transistor chip leaks more than one watt—not good for battery life in an "always-on" portable device. As transistors get smaller, leakage power increases. At 65 nm, the same 200-million-transistor chip leaks more than 10 watts.

Imagine the broad spectrum of applications, from smoke detectors to supercomputers. Some applications, such as wristwatches, want to run for years on a cheap battery. Some, such as weather modeling, want all the computational horsepower they can get and don't mind using lots of power to get it. Other applications, such as cell phones and digital cameras, need to balance power conservation and computational ability. This huge range of applications consumes billions of microprocessors each year. Curves like those in fig. 8 can represent each of the applications to show us the value transistor for that application. The point is that *each process generation carves off a slice of applications for which it makes the value transistor*. (Each process generation now retains a market segment.) Applications for which the semiconductor process has attained the value transistor now possess the right transistor for what they do and won't help pay for further process development. More process generations slice wedges off the application space until the applications that remain cannot afford to pay for the next process advance.

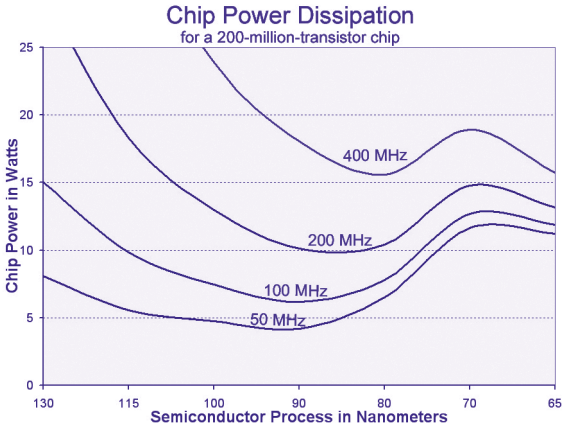

Fig. 8. Power for a 200-million-transistor chip at several frequencies. At 50 MHz, the value transistor occurs in a 90-nm process. At 400 MHz, the value transistor occurs in an 80-nm process.

sor brain improved its performance. Performance improvements in the PC's microprocessor came primarily from Moore's-law improvements in the semiconductor process. Until it reached a point of diminishing returns in the 1990s, the PC's system performance grew at a rate close to the Moore's-law rate of improvement in its microprocessor.

Demand for performance established the market. The PC's early adopters had high expectations, and their expectations grew with time. But, as the PC's market grew, its customer base grew to include late adopters whose expectations started lower and grew more slowly. The demand for performance thus spread along a range and grew at a rate slower than the supply. Eventually, the faster-rising supply curve crossed the lower boundary of the spread-out demand curve. That meant that the PC's performance was higher than the demand, for one segment of the market. Customers began buying cheaper "value" PCs that offered enough performance.

Moore's law supplied performance, and the PC eventually overshot the demand for much of its market. Moore's law supplies performance in transistors too. Just as the PC's performance overshot some of its market, a transistor's performance can overshoot some of its applications.

A battle raged for decades in the engineering community over whether assembly-language (computer specific) programming was better than high-level-language (computer independent) programming. Arguments focused on performance and on program size. Advocates for assembly-language programming won those battles, but high-level languages won the war. The high-level languages won because the real issue was programmer productivity. High-level languages forfeited efficiency in algorithms and in program size to increase the productivity of the critical resource: the programmer.

Assembly-language programmers thought their battle was about performance; it was about productivity. Workstation makers thought their battle was about performance; it was about unit volume. ASIC advocates fight PLD advocates about performance; the battle is about what's good enough for the application. Moore's-law advocates think smaller, faster transistors drive the semiconductor market. Smaller, faster transistors are better only up to a point, and Moore's law is an enabler, not a driver.

The transistor's quest for performance is just what the PC needed. But the world is splitting into tethered devices, such as PCs and printers that plug into power outlets, and untethered devices, such as cell phones and personal digital assistants that carry their power sources. As the market for value PCs grows, engineering resources will be reallocated to more profitable untethered applications. The transistors supporting untethered applications need to balance cost, performance, and battery life. Chips for untethered applications need the right transistor for the job; that isn't necessarily the smallest, fastest transistor.

Just like the assembly-language programming advocates, today's engineering community focuses on the transistor's size and speed. But the market is changing, so the transistor's quest for performance is now misdirected.

Lessons

There are a lot of "what ifs" possible with a chip manufacturing model. This sequence wasn't meant to be definitive, but to be qualitatively correct. I believe that the situation is as follows. Twenty years of building PCs for performance has brought the industry to the belief that Moore's law drives the semiconductor industry. It doesn't. Moore's law is the fantastic *supply law* that is the industry's hallmark. Moore's law is an enabler—it says what is possible. It says nothing about *demand*. Moore's law has been the response to demand.

I don't work in a foundry, so my model might be off in detail, but it illustrates important lessons for the industry. Fixed costs for the plant and its equipment and for process development are amortized across all wafer production in the plant. Fixed costs for masks are amortized across the wafers in the production run for a specific chip. All of the fixed costs are increasing.

Escalating fixed costs have two important consequences. First, a chip's equal-cost crossover point between an old process and a new one is moving toward *much higher* production runs. Unless you are building millions of systems, your chips will be cheaper in the *old* process. Second, the advantage of moving from an old process to a new process decreases in both absolute dollars and as a percentage of the cost for each new process generation. In older processes, a redesign might have saved 50%, while future processes promise only 10%.

Because its fixed costs have been amortized, a 250-nm process builds 25-million-transistor chips that are cheaper than 25-million-transistor chips built at 180 nm or at 130 nm. Once its fixed costs are amortized, the 180-nm process will build the cheapest 50-million-transistor chips (Moore's law). Any application that isn't performance-limited in 180 nm and doesn't need more than 50-million transistors (a *lot* of applications) will be cheaper in the fully amortized 180-nm process than in any process with smaller transistors. That's a huge chunk of the application space that is doubling with each process generation.

Advocates with leading-edge performance and capability requirements make the case for following

Moore's law to the next smaller, faster transistor. Leading-edge applications need the performance or capability and are willing to pay a premium to get it, but these applications aren't the bulk of the market. Further, with each process generation fewer applications remain to pay the next generation's escalating costs. Intel leads the charge, betting that demand will escalate with fab cost.

The PC dominated the market for twenty years, but requirements are changing. As the PC market moves from performance to value, engineering emphasis will move to untethered applications. Untethered applications change the transistor's requirements from absolute performance to a balance of cost, power conservation, and performance. For untethered applications the right transistor for the job, the value transistor, may not come from a leading-edge process. Big transistors burn more active power; small transistors leak more.

Escalating costs have to be amortized over a fixed interval, which is causing a problem. All the while it's leaving a wake of processes that are good enough (their applications will never migrate).

Advances can change the rules, right? Maskless techniques, such as e-beam lithography, could drop mask cost to zero. Double- and triple-gate transistors can substantially reduce leakage power, which alters the position of the value transistor. Silicon-on-insulator, strained silicon, and exotic materials may change tradeoff points. The move from 200-mm wafers to 300-mm wafers affects the economics of what process is appropriate and what production runs are cost effective. So, yes, advances can change the rules. But advances take time, and the cost to change from an old process to a new process increases with time while the payoff decreases. So, for a large segment of applications, the transistors have become *good enough*.

Moore's law is an enabler, not a driver. Since the transistors are now good enough for a wide range of applications, the semiconductor industry can be healthy even if equipment makers can't sell more advanced equipment. The emergence of the value transistor invalidates "leading indicators" that assume transistors must get smaller to get better and, therefore, require the newest equipment to sustain industry growth. Up to now, the transistor wasn't good enough, and all applications shared the cost of process advances; in the future, huge segments of the market will reach their value transistor and will no longer share the cost of process advances.

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Chartered	Foundry	Excellent	As a demand-driven supplier, Chartered is in a good position to produce value transistors.
GSMC, SMIC	Foundry	Excellent	Grace Semiconductor Manufacturing Corp. and Semiconductor Manufacturing International Corp. are Chinese foundries. With huge and growing domestic and export markets, these companies are perfectly positioned to exploit growth in production of value transistors.
National Semiconductor	Integrated	Good	National has its own fabs but already produces a range of products that are more cost oriented than performance oriented. National's fabs can profitably build value transistors.
TSMC, UMC	Foundry	Good	Foundry prospects would be excellent except that as these companies move production to China, they are hamstrung by Taiwanese export restrictions to a 250-nm process.
Altera, Xilinx	Fabless	OK	Altera and Xilinx are among a diminishing number of companies left to share the cost of advanced process development (by paying premiums to TSMC and UMC for leading-edge chip production).
Motorola, Samsung, Texas Instruments	Integrated	OK	Motorola, Samsung, and TI invest in leading-edge processes but have business models more closely tied to demand than does Intel.
Applied Materials	Systems	Struggle	Applied Materials is caught in the momentum of producing high-end semiconductor processing equipment. The market is moving to value transistors, which require cost-reduced, low-end production equipment.
IBM	Integrated	Struggle	IBM invests heavily in leading-edge process development, and it charges a premium for its chips. Demand for leading-edge processes will fall as more applications are satisfied with already available transistors.
Intel	Integrated	Struggle	Intel's process development supports its high-end microprocessor business model. As demand for transistors shifts from performance to value, Intel will struggle.

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	11/30/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	14.53	8.32 - 27.59	5.56B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	30.69	17.88 - 48.84	11.2B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£0.34	£0.26	£0.20 - £0.59	£0.76M
ARM Limited (ARMHY***)	Microprocessor and Systems-On-Chip Cores	11/26/01	16.59	3.28	1.87 - 17.94	1.1B
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	6.11	2.23 - 6.53	187.4M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	6.44	4.63 - 30.36	1.0B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	8.64	3.60 - 26.20	1.1B
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
Energy Conversion Devices (ENER)	Ovonic Unified Memory	6/18/02	27.69	12.16	7.21 - 25.73	266.3M
Flextronics International (FLEX)	Contract Manufacturing	8/6/02	7.68	11.01	5.47 - 29.99	5.7B
Foveon (none*)	CMOS Imaging Chips	6/18/02				
Legend Group Limited (LGHLY.PK)	PCs and Consumer Electronics	8/6/02	6.63	7.15	N/A	N/A
Microvision (MVIS)	MEMS-based Micro Displays, Nomad Head-Worn Display, Scanners	6/18/02	6.80	5.52	2.64 - 15.50	83.7M
National Semiconductor (NSM)	Geode x86 Microcontrollers, Consumer Orientation, 51% Ownership of Foveon	6/18/02	32.30	20.30	9.95- 37.30	3.7B
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 ††	9.25	5.31 - 19.08	34.2B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	1.52	0.74 - 4.47	206.4M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	4.54	2.93 - 10.02	13.9B
VIA Technologies (2388.TW)	x86 Microprocessors for "Value" PCs	6/15/02	78.00	43.70	39.00 - 127.87	N/A
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	6.06	2.03 - 20.14	479.6M
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	24.64	13.50 - 47.16	8.3B

† 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.