Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 11 | The Investor's Guide to Breakthrough Micro Devices | November 2002

A Bridge Too Far

The big foundries, Taiwan Semiconductor Manufacturing Corporation and United Microelectronics Corporation, have curtailed capital expenditures and are pushing out adoption dates for 300-mm (diameter) wafer production. The biggest integrated device manufacturers (IDMs), IBM, Intel, Samsung, and Texas Instruments, continue to spend on process development and are moving to 300-mm wafers. Conventional wisdom says that as the economy recovers, IDMs will be prepared for the upturn. And foundries, with lagging semiconductor processes that produce slower chips at higher costs, will lose market share.

In the early days, foundries lagged IDMs by at least a couple of process generations. Then they caught up with the IDMs—introducing large wafers and leading-edge processes right along with the major IDMs. Now, foundries seem about to fall behind. What is going on?

Moore's law is the rate of semiconductor manufacturing improvement—the number of transistors in a fixed area doubles every eighteen months. Big chips are more capable; fixed-size functions fit on smaller chips. The magic of Moore's-law progress comes from three areas. Most important is shrinking transistor size. The second contributor is increasing chip and wafer size. Third is better circuit design. Chips get faster and cheaper with manufacturing process improvements. If you find the current chips lacking, wait a generation or two and they'll have what you need. If you find the current chips capable, but they're too expensive, wait a generation or two and they'll be cheap enough. That's the way it's been for over thirty years.

100 nanometers = 0.1 microns = 0.0001 mm
Example: 130 nm = 0.13 microns
Bacterium = 1,000 nanometers wide

But there's an interesting chart at TSMC's web site (www.tsmc.com/english/technology/t0203.htm) that isn't easy to explain. Fig. 1 is a version of TSMC's chart with its interesting features. The unit of measure for semiconductor manufacturers is "wafer starts." Chip size and transistor size vary with product and process; wafer size stays constant for years, so production capacity is measured by the number of wafers the plant processes per month.

Fig. 1 shows the percent of the foundry's wafer starts, by semiconductor process, plotted against time. Before the 1996 introduction of the 350-nm process, 100% of the foundry's wafers were at 500 nm or larger. By the beginning of 1997, more than 30% of wafer starts were at 350 nm. In 1997, the foundry offered a 250-nm process, but by the end of the year, fewer than 20% of its wafer starts were at 250 nm. This is an interesting trend: with each new generation of semiconductor process, the adoption rate is falling.

Moore's law says there's incentive to move to a more advanced process. Chips get smaller and faster, they use less power, and they are cheaper. Let's say you are making chips in a 500-nm process and have an opportunity to move to a 350-nm process. We've talked about the speed and power advantages, so

In This Issue:

Even big things—manufacturing, personal computers, transistors, white-collar labor—have a life cycle. It's a characteristic of big things that when we are in the middle of them we can't imagine how they could ever plateau. Afterwards, we wonder how we could not have seen it. PCs are not coming to an end, but they are becoming good enough. And the same thing is happening to the transistor. The implication is that new semiconductor processes are not worth doing.

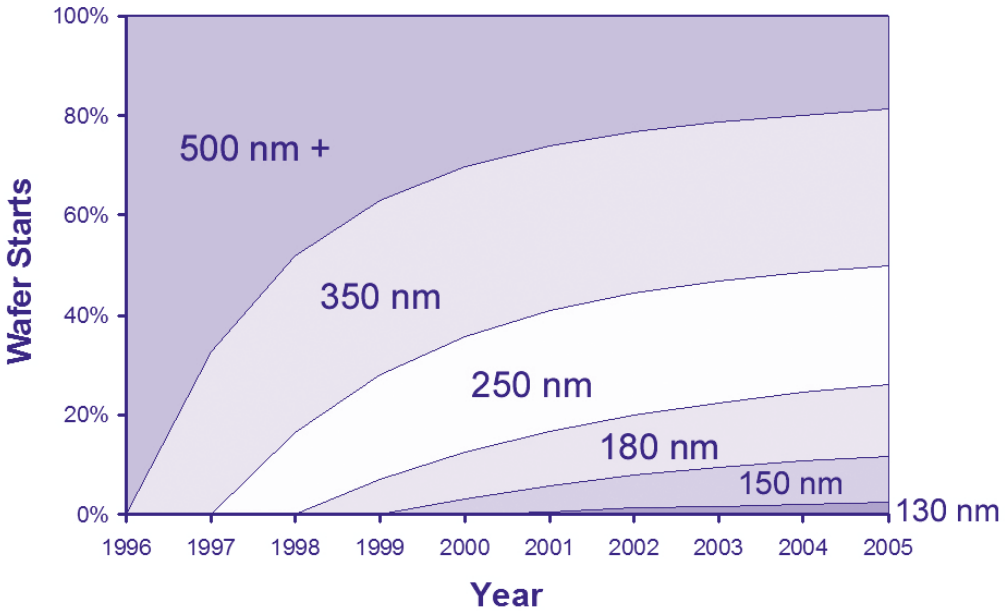

Fig. 1. As the foundry moves to new process geometries, *adoption rates decline*.

what about cost? To a first-order estimate, the same wafer size should hold twice as many chips in a 350-nm process as it does in the 500-nm process. It costs about $500 to process a 200-mm wafer. This cost, the cost to *operate* the wafer-processing equipment, is independent of whether the wafer's patterns are at 350 nm or at 500 nm. If you sell a wafer's worth of 500-nm chips for $1,000, then $500 is profit. If you sell them at the same price, a wafer's worth of 350-nm chips will fetch $2,000, with $1,500 of profit. You might charge more for the faster, lower-power 350-nm chips, making margins even higher.

Fig. 2 shows how smaller chips pack better on the wafer. The chip edges have to line up in both directions so the processed wafer can be cut into chips. While a 50% shrink in each dimension should allow 4 times as many chips, the actual number is more like 4.5 because the smaller chips pack better on the round wafer.

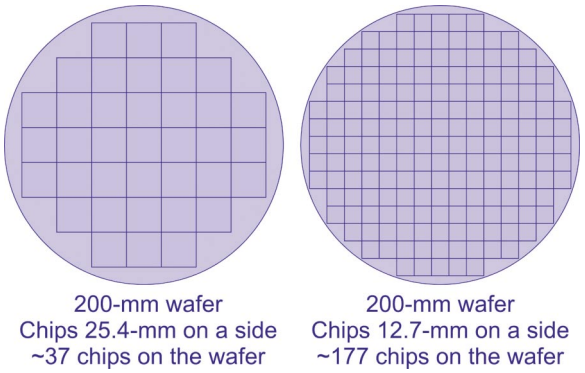

Fig. 2. Smaller chips pack better on a round wafer. Halving the chip dimensions means more than four times as many chips on the wafer.

There's more incentive than just how many chips *fit* on the wafer. Defects reduce the wafer's yield of good chips. Yield is the percent of good chips per wafer. Fig. 5 illustrates the effect of ten random defects on the wafer's yield, for two chip sizes. For this example, moving from a half-inch chip to a quarter-inch chip *improves the yield* from 75% to 94%. Therefore, instead of 4.5 times as many chips, the same-size wafer yields almost 6 times as many *good chips*. Financial incentives for moving to smaller geometries seem compelling.

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.

How Transistors Shrink

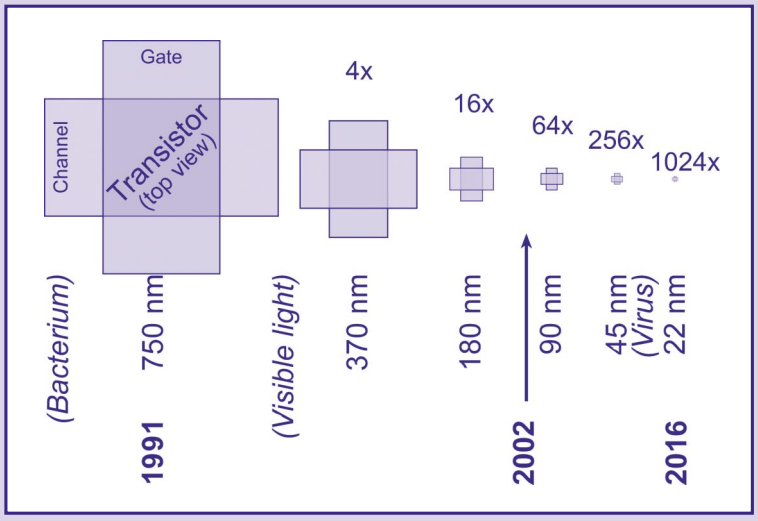

Fig. 3 illustrates both the magic and the astounding challenges of shrinking CMOS transistors. The simple intersection between a "channel" and a "gate" forms a transistor. The channel is on the bottom and is like a highway in silicon. The gate crosses the highway at a right angle. Voltage (electric force) on the gate controls whether electrical current flows in the channel.

Fig. 3. Moore's law progress shrinks transistors.

Fig. 4 shows the parts of a typical transistor, called a field-effect transistor. Positive voltage on the gate causes electrons to collect at the border between the pure-silicon channel and the insulator. These electrons form a bridge so that electricity can flow through the channel. Big transistors use more power and are slower than small transistors because the bigger gate employs more electrons to turn the transistor off and on. The insulator between the gate and the channel keeps electrons from leaking between them. The thicker the insulator, the higher the voltage must be on the gate to turn the transistor on. That's the basics of transistor operation. Now, for the interesting effects of scaling.

Fig. 4. Side view of a typical transistor. Voltage on the gate collects electrons in the channel to allow current to flow under the gate. The gate acts as a valve to control the flow of electrical current in the channel.

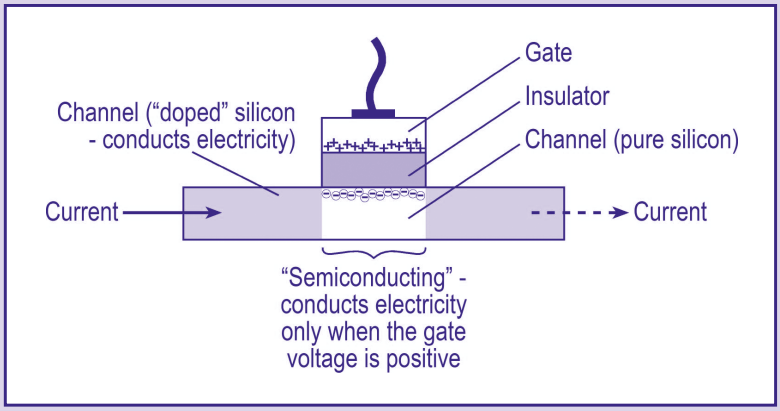

As the process geometry shrinks, the transistor shrinks in both directions: the channel narrows and the gate narrows (see fig. 3). When dimensions shrink by half, four transistors fit in the area formerly occupied by one transistor. By 1991, manufacturers built chips with 750-nm features—already smaller than a bacterium. Three years later, in 1994, features were 370 nm—*smaller than the smallest visible light waves*. Transistors with 370-nm features cannot be seen with an optical microscope. Today's leading-edge process builds sixteen transistors in the area occupied by that 1994 transistor that was too small to be seen with visible light. By 2016, transistors will be smaller than a *virus*—sixteen of these year-2016 transistors will fit in the area of today's leading-edge transistor.

As the transistor gets smaller, it deals with fewer electrons, so it turns on and off faster and it uses less power. As the insulator between the channel and the gate gets thinner, the transistor can operate with lower voltage and with less power. But as the gate that controls the channel gets narrower, electrons leak across the pure-silicon barrier (like a dripping faucet). And as the insulator between the gate and the channel gets thinner, it can leak and may allow current to flow where current should not be present (like a faucet with water leaking out of the handle).

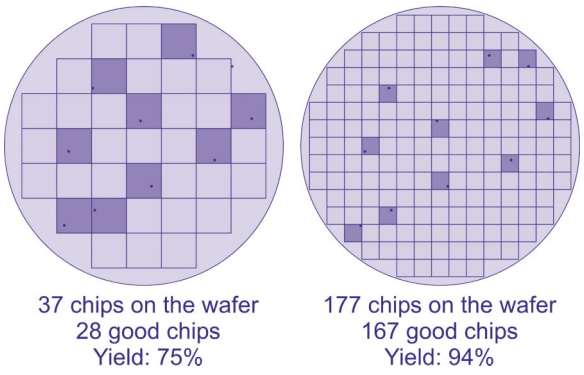

Fig. 5. The yield can be substantially higher for smaller chips. Halving the chip dimensions yields six times as many good chips.

Fig. 6 shows the advantage when the foundry moves from 200-mm diameter wafers to 300-mm diameter wafers. Just based on area gain, we expect 2.25 times as many chips on the larger wafer. The number will be slightly higher because chips pack better on the larger wafer. Intel, for example, gets 201 of its 134-square-mm Pentium 4s on a 200-mm wafer and 482 Pentium 4s on a 300-mm wafer. It's about 2.4 times the number of chips, but the cost to process a 300-mm wafer should be only 20% higher than the cost to process a 200-mm wafer.

It seems financially compelling to move to smaller processing geometries and to move from 200-mm wafers to 300-mm wafers. But, as fig. 1 shows, at least for the foundries, the move isn't happening.

Wafer starts for old semiconductor processes don't fade to nothing. Instead, they decline and then stabilize for years as a percent of the foundry's wafer starts. Why might this be so? And what does it mean for Dynamic Silicon companies?

The answer to "why chip designs don't move to new processes" is threefold: start-up costs, hidden costs, and physical limits.

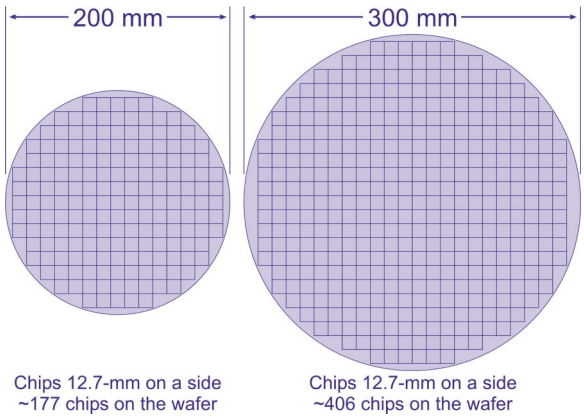

Fig. 6. Larger wafers hold more than twice as many chips, but cost only 20% more to process.

Start-up costs

Theoretically, it costs about the same to process a 200-mm wafer whether the lines on the wafer are 180-nm wide or are 130-nm wide. This cost (wafer processing) is a variable cost. Practically, however, the foundry's 180-nm plant will be two years old, while its $2.5-billion 130-nm plant will be new. This $2.5 billion is the fixed cost of the building and the cost of the processing equipment with $500 to $600 million in process development cost (the cost to develop design rules for the 130-nm process). Wafers running the 130-nm process in the new plant have to pay their share of fixed costs. The fixed costs of the old process have been amortized over two years of production. The foundry can charge less for work in the depreciated 180-nm plant.

Hidden costs

You can't just call the foundry and say: "Move the Umptyfritz Controller production from the 350-nm plant to the 180-nm plant." Your engineering teams developed the controller for a specific 350-nm process. If you want to produce it at 180 nm, your engineering teams have work to do. They have to build the Umptyfritz Controller for the 180-nm process. For most engineering managers, the choice comes down to allocation of precious engineering talent. Do you want your engineers cost-reducing an old product or do you want them working on the next-generation product?

If your engineers cost-reduce the product, one cost will be a new mask set. Mask sets are expensive and they are getting even more expensive. As a rule of thumb, each new process generation doubles the cost of the mask set. For a 130-nm process, mask costs can be $1 million. At 90 nm, masks could cost $2 million. By 2010, a mask set could cost $10 million! If the chip doesn't work, you'll have to buy more than one mask set. Mask costs have to be amortized over the production run. If the chip you build goes into a system with expected lifetime sales of 100,000 units, then the $1 million mask cost adds $10 to the cost of each chip. Amortized engineering cost adds more. Continued production of the old chip may be cheaper.

The chip design tools the engineers used for the 350-nm process don't work for the 180-nm process. The engineers need new tools.

Four hidden costs in moving a chip from an old process to a new one are time, engineers, masks, and design tools. Moving from an old process to a new one will get simpler as chip design descriptions get "softer" (*Dynamic Silicon*, Vol. 1, No. 6), but it still costs time, talent, and money. (Soft descriptions are parameterized recipes that are independent of the manufacturing process.)

Physical limits

If the chip doesn't get smaller, there's no cost advantage in moving to a more-advanced process. Making the *circuits* smaller may not make the *chip* smaller (*Dynamic Silicon*, Vol. 1, No. 3). This is because the chip's connections to the outside world are through wires attached to "bonding pads" on the chip. Bonding pads can't shrink below the area that an automated bonding machine can hit.

There are two lower bounds on a chip's size. One is the area occupied by the circuits and the other is the minimum area enclosed by the bonding pads. If the bonding pads determine the minimum chip size, the chip is "pad limited" (fig. 7). If the chip is pad limited in the 350-nm process, there's no cost advantage in moving to the 250-nm process.

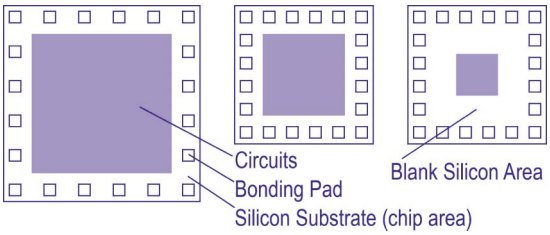

Fig. 7. Pad limited: as the semiconductor process improves, the chip and the circuits shrink, and the bonding pads get closer together (middle chip) until the bonding pads limit the size of the chip (right chip).

Microcontrollers for consumer appliances are good examples of pad-limited designs. A hair dryer, a microwave oven, or a blender can do all it needs to do with an eight-bit microcontroller that would be pad limited in the most advanced process.

A pad-limited design in one semiconductor process has twice the number of transistors (Moore's law) as a pad-limited design in the previous process generation. Each process increment, therefore, enables a new wedge of applications that will never need a better process. That takes a bite out of the demand for the next-generation process.

It's supply and demand again. Moore's law is supplying transistors. It doubles the supply of transistors every eighteen months. The supply of transistors exceeds the demand for many applications. As the number of transistors in even pad-limited chips gets large, chips cross the demand curve for more applications.

I've said that Moore's-law progress has left a huge wake of enabled but unexploited applications (*Dynamic Silicon*, Vol. 2, No. 1). The persistent residual percentages of wafer starts in old processes are evidence that it's true. The *total* number of wafer starts grows each year. If the 350-nm process is still 20% of wafer starts six years after its introduction, then the demand for the 350-nm process is growing at the rate that all wafer starts are growing.

The minimum unit that a foundry processes is a "boat" of twenty-five wafers. The lot size, or number of chips in a boatload, for the big chips on 200-mm wafers in fig. 2 is about 1,000. Shrinking the geometry by half brings the lot size to about 5,000. Moving to 300-mm wafers and shrinking the geometry by half again increases the lot size to about *50,000*. If you don't need chips in 50,000-unit quantities, you may not need an advanced process and 300-mm wafers.

Also, if the bill of materials for your system is $1,500, cost-reducing a two-dollar microcontroller won't be your top priority.

Machine tools and mass production

Eli Whitney invented interchangeable parts in the same sense that Columbus discovered America or that Intel invented the microprocessor (the real stories are more complicated). In 1798, Eli got a contract to build 10,000 rifles for the army in two years. Rather than building parts, Eli Whitney built *machines* that built parts. This was a new idea. It took ten years to get it right, and it began the machine-tool industry. The introduction of interchangeable parts spawned the Industrial Revolution in Great Britain and in the United States, but with different objectives. Great Britain had plenty of skilled labor, so the Brits' machine tools increased component precision. In Great Britain, machine tools made skilled workers better. The United States was short of skilled labor; machine tools, worked by relatively unskilled laborers, could turn out large quantities of identical components. In the United States, machine tools amplified the workforce. Machine tools raised the level of abstraction from component *fabrication* to machine-tool *operation*. Large numbers of unskilled workers, making components on machine tools, leverage the work of one skilled machine-tool maker.

Interchangeable parts enabled the assembly line; mass production followed. Mass production led to huge increases in worker productivity. Wages rose, production increased, and costs dropped, stimulating the Industrial Revolution's economic growth.

Vertical integration

Production automation continued when Ford opened its River Rouge plant in 1927 building the Model A Ford. This 2,000-acre plant was the ultimate in vertical integration. It had its own glass, tire, and engine manufacturing, docks, blast furnaces, steel mills,

foundries, rolling mills, metal stamping, and even its own power plant supplying electricity and steam. Iron ore and other raw materials, received daily at one end of the plant, rolled out the other end twenty-eight hours later in a finished automobile. General Motors became similarly vertically integrated.

Dealing with suppliers carries a "transaction cost." Prices, delivery, component specifications, quality, quantity, and schedules have to be negotiated. Negotiating with suppliers isn't building products; it's overhead. Further, there's no real control over the suppliers; suppliers may or may not deliver. Therefore, buy the supplier. Corporate fiat then replaces negotiated transaction costs. But now the company incurs transaction costs dealing with the suppliers to the original supplier. Instead of negotiating with the supplier of an electrical alternator, the company now negotiates with the suppliers of castings, wire, diodes, bearings, bushings, insulators, and pulleys. Therefore, buy suppliers until you reach the suppliers of raw materials.

That's what the automotive industry did for decades. Now the automobile industry is fragmenting horizontally. Vertical integration has its advantages, but it has corresponding disadvantages. Internal suppliers have a captive market and are not subject to competitive pressures. Costs are difficult to estimate for internal suppliers. There are still transaction costs, but they are internal so they're harder to see and they're harder to measure.

Standards motivate horizontal fragmentation. Should GM's Delco Electronics division continue to supply only GM, or should it be spun out to Delphi? Standards make electronic subsystems more interchangeable. GM transferred Delco Electronics to Delphi in 1997. An independent Delco might achieve economies of scale unattainable as a captive supplier. It will have to compete with Infineon and with other makers of automotive electronics even for GM's business, putting pressure on the company to be efficient. It can serve customers outside the automotive business, increasing scale.

The automotive business began with hand-crafted designs, and it shifted to mass production with the introduction of interchangeable parts. The lack of suppliers and of standards led to vertical integration. But the emergence of reliable suppliers and of standards encouraged horizontal fragmentation to maintain economies of scale.

Lessons

What does this have to do with silicon? Plenty. Mass production, interchangeable parts, economies of scale, vertical integration, and coming horizontal fragmentation all apply to the semiconductor business.

The invention of the integrated circuit (chip) began the business of building custom chips. The first companies building chips were vertically integrated. They grew silicon crystals, sliced them into wafers, designed the circuits, processed the wafers, packaged the chips, and sold the final product. They even built their own semiconductor processing equipment.

Integrated-circuit macros (e.g., "TTL" parts) were the semiconductor industry's first interchangeable parts, and they spawned a booming circuit-board industry that fueled the growth of electronics applications.

The microprocessor, which Intel introduced in 1971, raised the level of abstraction for designing electronic systems. Lowering the skill needed for electronic design raised the pool of design engineers, and it accelerated the growth of the industry. Unlike application-specific integrated circuits, which serve only the limited quantities required for a specific application, general-purpose microprocessors are mass produced to be used in many applications. Economies of scale reduced the microprocessor's cost, and further broadened its application.

As the industry grew, it became cost effective to put ancillary functions on the chip with the microprocessor, creating the microcontroller. The microcontroller with its integrated peripherals reduces the chip count and the cost for a narrow range of applications. Microcontroller varieties proliferated, fragmenting the application space.

The rise of standards is doing to the semiconductor industry what it has been doing in the automotive business. Semiconductor equipment is now standard across the industry; no IDM builds its own processing equipment. It's more efficient for a few suppliers to grow silicon crystals and to supply blank wafers to the industry. The physical and logical interfaces of electronic components become standard across the industry, making their design by a captive engineering group inside a single IDM inefficient.

The semiconductor industry's boom and bust cycles encourage fragmentation. IDMs cannot afford the production capacity to sustain boom cycles because this capacity will be idle during bust cycles. But if the IDM doesn't have the capacity for the boom, it loses market share to competitors. Better to build capacity for average demand and contract with foundries for excess production. Contracting with foundries is feasible only if the processing is compatible. Compatible processing has the advantage of sharing the process-development cost, but it forfeits a primary advantage of owning a fab—a custom process.

When the PC came out, it didn't perform well

enough to satisfy anyone. It wasn't even close. PC makers could sell new, higher-performing models to anyone. But the supply of performance eventually exceeded the demands of low-end users. PC makers began to lose former high-end sales because customers bought "value" PCs instead. Now the low end has become the high-enough end for many people. The same thing is happening with the transistor. It's why the semiconductor process adoption lines are falling over in fig. 1. When transistors were big, everyone wanted better ones. As the transistor shrinks, more applications are satisfied with available "value" transistors. That leaves fewer applications that are willing to *pay* for the next process improvement.

Moore's law says what's *possible when demand is assumed*. When the transistor wasn't good enough, every application was willing to share the cost of Moore's-law semiconductor process advances. But, combine weakening demand with the escalating cost for process development, mask sets, and production equipment, and eventually there won't be enough demand to support moving to the next process. It's a race to see whether Moore's law or the customer base to support its advance runs out first. I'm betting the base demand out first. The transistor will reach its *practical limit* for the vast majority of applications before it reaches physical limits. The foundries aren't falling behind the IDMs in process development or in moving to 300-mm wafers; they are merely reflecting their customers' demand.

In *A Bridge Too Far*, Cornelius Ryan tells the huge cost of attempting to reach an overly ambitious objective (Operation Market Garden to capture the bridge at Arnhem [Holland] in World War II). The reason that the foundries seem to be losing ground to the IDMs in semiconductor process development may be that the IDMs have gone "a bridge too far." The IDMs *tie business models to semiconductor process development*. Intel, for example, invests revenue from its leading-edge microprocessors in process development to build the next leading-edge microprocessor. The IDMs build chips in leading-edge processes and hope the customers will be there. They will know they have gone "a process too far" when demand falls. By contrast, customer demand drives the mix of semiconductor processes at the foundries. As fig. 1 shows, adoption rates are falling for newer processes. The foundries are in the right place at the right time; the IDMs will fall with declining advanced-process adoption rates.

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Altera, Xilinx	Fabless	Excellent	Programmable logic components are the ultimate in interchangeable parts. These chips can use any process from high margin, high performance to obsolete. Good margins at the high end subsidize development of core IP libraries.
Cypress Microsystems, Triscend	Fabless	Excellent	Great position to consolidate microcontroller proliferation. Post-production customization means fewer chip types and higher production volumes, which reduces costs.
GSMC, SMIC	Foundry	Excellent	Grace Semiconductor Manufacturing Corp. and Semiconductor Manufacturing International Corp. are Chinese foundries. They are beginning with 250-nm process and will soon move to 180 and 130, unhampered by external investment rules.
ARC, ARM, Tensilica	Fabless	Good	Microprocessor cores are the building blocks for the next generation of system-on-chip designs. Soft IP cores are portable across foundries and processes.
National Semiconductor	Integrated	Good	Diverse product line offsets disadvantages of owning semiconductor production plants. It has advantages in producing x86-based microcontrollers for consumer applications.
Transmeta, VIA Technologies	Fabless	Good	Well positioned for "value" applications of x86. Their opportunity would improve if they offered x86 cores.
TSMC, UMC	Foundry	Good	Foundry prospects would be excellent except that as these companies move production to China, they are hamstrung by Tiwanese export restrictions to 250-nm process.
Analog Devices	Integrated	OK	Diverse product line offsets disadvantages of owning semiconductor production plants.
IBM	Integrated	OK	Internal contention system offsets some disadvantages of vertical integration, making the horizontal fragmentation less difficult than it will be for companies with a monolithic culture.
Samsung, Texas Instruments	Integrated	OK	Diverse product line offsets disadvantages of owning semiconductor production plants.
Intel	Integrated	Struggle	Intel's process development supports its high-end microprocessor business model. As demand for PCs shifts from performance to value, Intel will struggle.
Motorola, STMicroelectronics	Integrated	Struggle	Circumstances that built vertically integrated semiconductor companies have changed to favor horizontal fragmentation. Vertically integrated companies will struggle with the transition.

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	10/31/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	11.72	8.32 - 27.59	4.47B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	26.80	17.88 - 48.84	9.8B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£0.34	£0.26	£0.20 - £0.64	£0.76M
ARM Limited (ARMHY***)	Microprocessor and Systems-On-Chip Cores	11/26/01	16.59	2.73	1.87 - 19.20	917.1M
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	4.96	2.23 - 6.79	152.1M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	5.00	4.74 - 30.36	693.0M
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	5.62	3.60 - 26.20	695.4M
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
Energy Conversion Devices (ENER)	Ovonic Unified Memory	6/18/02	27.69	9.92	7.21 - 25.73	217.2M
Flextronics International (FLEX)	Contract Manufacturing	8/6/02	7.68	8.36	5.47 - 29.99	4.32B
Foveon (none*)	CMOS Imaging Chips	6/18/02				
Legend Group Limited (LGHLY.PK)	PCs and Consumer Electronics	8/6/02	6.63	6.80	N/A	N/A
Microvision (MVIS)	MEMS-based Micro Displays, Nomad Head-Worn Display, Scanners	6/18/02	6.80	4.60	2.64 - 16.00	66.6M
National Semiconductor (NSM)	Geode x86 Microcontrollers, Consumer Orientation, 51% Ownership of Foveon	6/18/02	32.30	13.28	9.95- 37.30	2.4B
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 ††	7.82	5.31 - 19.08	28.9B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	0.91	0.74 - 4.47	121.7M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	4.15	2.93 - 10.02	11.1
VIA Technologies (2388.TW)	x86 Microprocessors for "Value" PCs	6/15/02	78.00	47.70	39.00 - 127.87	N/A
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	3.62	2.03 - 20.14	286.5M
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	18.99	13.50 - 47.16	6.4B

† 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.