Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 4 | **The Investor's Guide to Breakthrough Micro Devices** | **April 2002**

Semiconductor Progress: Trends and Consequences

In fifteen months of Dynamic Silicon, I've covered Moore's law, shrinking computers, the horizontal fragmentation of chip manufacturing, the evolution of digital design, component proliferation and consolidation, and a bunch of other trends. This month I'm tying these ideas into a story. Yes, the Big Picture. Interdependencies and subtleties make the big picture elusive. Inventions and new products create opportunity and are the seeds of change, but consequences may lag change by years. Before I start, however, I cover recurring themes. These themes—changing supply and demand, proliferation and consolidation, vertical-to-horizontal fragmentation, and scarcity and abundance—are the forces that drive change.

Timeless themes

Changing supply and demand. Laws of supply and demand drive markets. The supply can be microprocessors, automobiles, personal computers, transistors, gigahertz, capacity—anything that can be measured. Demand is the market's "pull" for what is being supplied. *Demand is difficult to measure because you can't see it directly*. Early PCs, for example, weren't fast enough. As manufacturers supplied more speed, users bought new PCs. Then, as the market grew, users' expectations rose, but expectations also widened. Now, the PC's rapidly rising performance exceeds demand for large parts of the market. The changing relationship between supply and demand precipitates change.

Proliferation and consolidation. As a new market grows, suppliers proliferate, then consolidate. The early days of automobiles and hard disks, for example, spawned dozens of makers. As these markets matured, the makers consolidated.

Vertical-to-horizontal fragmentation. In developing markets, suppliers are vertically integrated; they build the system from top to bottom. In computer mainframes, for example, IBM built everything from components through systems and software. As the computer industry matured, the industry fragmented horizontally into semiconductor equipment makers, chip suppliers, system houses, and software and operating system suppliers.

Scarcity and abundance. Successful problem-solving spends the abundant resources to conserve the scarce resources. Computer programmers are scarce. The rising level of abstraction in programming is evidence of this shortage. Assembly-language programs use less memory and run faster than high-level-language programs, but high-level languages improve programmer productivity. The computer industry trades abundant memory and performance for programmer productivity.

The integrated circuit (IC): 1959

An IC is called that because it allows different electronic elements—transistors, resistors, capacitors, wires—on one silicon chip. The integrated circuit ignited the semiconductor industry's "magical cycle" of ever increasing function and ever lower cost.

Moore's law. Moore's law describes the magical cycle. The *International Technology Roadmap for Semiconductors* (ITRS, 2001 Edition) projects that by 2016 semiconductor line widths will be 22 nanometers—10 times the width of a DNA molecule. Today's high-end

Relative Sizes

	Width in nanometers
Dust mite	300,000
Grain of sand	100,000
Human hair	75,000
Human cell	10,000
Bacteria	1,000
Red light	600
Hard-disk track	400
Chip line width (2002)	100
Hard-disk bit	40
Virus	35
Chip line width (2016)	22
Protein molecule	12
DNA molecule	2.4
Silicon atom	0.23

1 Micron = 1,000 nanometers

chips contain hundreds of millions of transistors. Chips of 2016 could contain tens of billions of transistors.

Decreasing semiconductor geometries enable more transistors on a fixed-size chip. More transistors offer more capability, so new chips have applications that were inaccessible to the previous generation.

Decreasing geometries shrink the area of a chip with a fixed number of transistors. Smaller chips are cheaper. So they reach cost-sensitive applications that were out of reach.

The application areas enabled by semiconductor manufacturing improvements far exceed the ability of available engineering talent to exploit it. It's supply and demand. This time the supply is sufficiently cheap transistors to enable cost-effective applications. Demand is limited by the availability of engineers to design chips into new applications. Moore's law progress is leaving a huge swath of potential applications that will be unexploited long after semiconductor geometries bump against physical limits.

Like food preparation, making chips has two parts: the process (recipe and techniques) and the fabrication plant (kitchen and equipment). From today's 130-nanometer geometries (one-quarter the wavelength of light) to tomorrow's 22-nanometer geometries, each process generation costs more to develop. Process development costs $500 million today. Fab costs exceed $2 billion today and could reach $20 billion by 2016. The escalating cost of process development and of the fabrication plant drives manufacturers to share process development and fabs. This cooperation drives process standardization. The escalating sophistication of semiconductor processing encourages specialized-equipment makers such as Applied Materials, Lam Research, and Credence Systems, who sell new equipment to each manufacturer with each process generation. The drive to semiconductor process standardization favors foundries (for economies of scale) and fragments the industry horizontally.

DynamicSilicon

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

Dynamic Silicon is published monthly by Gilder Publishing, LLC. Editorial and business address: 291A Main Street, Great Barrington, MA 01230. Editorial inquiries can be sent to: bozo@gilder.com. To request permission to republish an article, call 413-644-2138. **To subscribe call 800-229-2573**, e-mail us at dynamicsilicon@gilder.com, or visit our website at www.dynamicsilicon.com

Masks, the sophisticated patterns defining the layers of a semiconductor, cost more with each process generation. About ten years ago, the leading-edge 0.8-micron process required twelve masks with a total cost of $18,000. Today's 0.13-micron process requires thirty masks with a total cost of $1,000,000. Rising mask costs discourage the short runs of limited-production chips, such as application-specific integrated circuits (ASICs), and encourage the production of generic chips. This has favored chips, such as microprocessors and digital signal processors (DSPs), that fit a broad range of applications. The beneficiaries have been microprocessor, microcontroller, and DSP manufacturers such as Analog Devices, Motorola, Intel, Texas Instruments, and others. This is the way it has worked for the last thirty years. Today, high-volume production favors programmable logic devices (PLDs)—chips that are all the same in manufacturing and are customized in the field. The beneficiaries are Actel, Altera, Atmel, Lattice, QuickLogic, Xilinx, and others.

As the semiconductor industry grew, "integrated device manufacturers" (companies with in-house chip fabrication) proliferated. But rapidly rising mask costs, process development costs, and equipment costs will force chip-manufacturing consolidation.

The design gap. Moore's law doubles the transistors on a chip every 18 months. The electronics market grows at 16-17% per year, averaged over its boom and bust cycles. The supply of transistors grows exponentially (as the product of these two trends). Who puts all those transistors into products? The supply of design engineers isn't growing exponentially. This situation, the supply of transistors growing faster than the supply of design capability, is called the "design crisis" or the "design gap."

To keep the gap from widening, designers must become more productive as fast as the supply of transistors grows. There are ways to close the design gap. One way reduces the skills required to consume transistors. There are about ten times as many embedded systems designers as chip designers. If the chip-design process is simplified to enable embedded systems designers to design chips, designer productivity increases eleven fold. For example, Celoxica builds software tools that turn software constructs into logic circuits for PLDs. Rather than having a design engineer convert an English specification into a logic circuit, tools convert a software program that mimics the circuit's behavior into a logic circuit.

Producing more copies of each design employs more transistors for the same design effort. Building a single design that works across many applications increases the designer's effectiveness. The microprocessor is a general-purpose design that substitutes for (displaces) a range of custom designs. Intel's 8051 general-purpose microcontroller is about 33,000 transistors and represents a few man-years of engineering design. In the twenty-five years it's been on the market, Intel has sold more than a billion units. Perhaps ten trillion transistors per engineering-design year—that's leverage!

A third method raises the level of abstraction in design so that the work of the designer represents more transistors. The designer whose design library consists of blocks of memory and logic, each of which may contain thousands or millions of transistors, uses transistors more rapidly than a designer whose repertoire consists of individual transistors. Raising the level of abstraction also simplifies design skills, enabling a larger pool of engineers.

In programming computers, we saw a rising level of abstraction first in the migration of applications from assembly language to high-level language. We see it in hardware design in the move from logic schematics to programming-like languages such as Verilog and VHDL. Verilog and VHDL require hardware design skills, so we'll see a further transition to C-language derivatives that enable software programmers to be hardware designers. The pool of programmers is perhaps ten times the number of Verilog and VHDL designers.

IC macros. After the invention of the integrated circuit, companies such as Texas Instruments, Motorola, Fairchild, and Signetics, exploited the abundance of transistors to offer chips as common building blocks or "IC macros"—the chip equivalent of Lego blocks. Building-block chips like, adders, multiplexers, and shift registers, raised the engineer's design productivity. Propelled by Moore's law improvements, families of IC macros grew in complexity and in variety.

The programmable logic device (PLD): 1966

Think of the PLD as a two-level device. One level holds logic elements and wires. The second level holds a "personalization" memory. Ones and zeroes in the personalization memory physically connect the logic elements and wires to make circuits.

Sven Whalstrom invented the PLD too early. In 1966, there weren't enough transistors on a chip to make his invention practical.

The microprocessor: 1971

Intel introduced the first commercial microprocessor in 1971. The first microprocessors weren't about computing and they weren't about performance. Early microprocessors provided the what-to-do-next logic, or what engineers call the "state sequencer" or "controller," that could be programmed to mimic some hardware functions and to control others. The general-purpose microprocessor saved engineers the effort of designing a state sequencer. The microprocessor, memory, and a few standard peripheral chips displaced a large number of IC-macro chips. When its performance was good enough, the microprocessor consolidated the IC-macros' functions into its software, displacing those chips.

The microprocessor simplified design by bringing the computer's programming model to embedded systems. Before microprocessors, the embedded systems designer implemented the structure and the procedure in hardware. Embedded systems design required proficiency in hardware logic. The microprocessor provided a state sequencer and a set of general-purpose computing resources (registers, arithmetic unit, logic unit, shifter, etc.). After microprocessors, the engineer supplied the procedure as software and the microprocessor supplied the structure (a computer). Programmers with some hardware-design skills could design embedded systems. The microprocessor traded transistors and efficiency for a larger pool of designers.

The microprocessor's programmed model raised the level of abstraction above the hardware skills needed to design with IC macros. The larger pool of designers both helped displace IC macros and increased the number of new applications. Being programmed, the microprocessor works with many applications, which means high-volume production. High volumes reduce chip cost, which further broadens the range of applications.

Microprocessors enabled more designers, had adequate performance for most applications, and reduced the number of chips in a system (which lowered system cost). Microprocessor shipments rose rapidly bringing fortune to Intel, Motorola, Texas Instruments, AMD, Fairchild, Signetics, Mostek, and a host of others.

Moore's law progress led to the proliferation of IC macros; their improving transistor density led to the microprocessor. The microprocessor consolidated IC macros and displaced them. But the microprocessor will also be a victim of Moore's law progress. As Moore's law enables more transistors on a chip, the microprocessor pulls its memory and peripherals onto the chip to become a microcontroller. Microcontroller variety increases as unique sets of on-chip peripherals customize the chip for a particular application. As we shall see, this proliferation of microcontrollers is vulnerable to consolidation.

The application-specific IC (ASIC): 1981

Microprocessors consolidated IC macros in low-performance applications. Where the microprocessor's performance wasn't good enough, engineers designed custom hardware using IC-macro chips. The founders of LSI Logic built their business displacing the separate IC-macro chips in high-performance designs by putting the equivalent logic on one custom chip: the application-specific integrated circuit.

The growth of ASICs benefited companies such as IBM, Lucent, LSI Logic, NEC, Fujitsu, Toshiba, and Hitachi. ASICs supply performance and capability (transistors for functions). The supply of performance and of capability follows Moore's law. The demand for performance and for capability grows more slowly than Moore's law. For ASICs, just as for PCs, the demand profile spreads out over time, with high-end developers demanding leading-edge performance and capability, and with late adopters who are satisfied with much less. The growing difference between available performance and capability and the demand for performance and capability leaves ASIC applications vulnerable to displacement.

The IBM PC: 1981

The introduction of the IBM PC split microprocessor design into embedded microprocessors, which emphasize flexibility and low cost, and into PC microprocessors, which emphasize performance.

PC microprocessors. For twenty years, the PC's need for performance drove the industry. The PC wasn't fast enough, so consumers purchased successive generations of higher-performing systems. The beneficiaries were the chip makers such as Intel, AMD, VIA Technologies, and Transmeta and system makers such as Dell, Compaq, and Gateway.

In PC microprocessors, more transistors per chip mean more performance. The demand for performance spreads out with time, creating a range of demand between high-end users and low-end users. For most of this range, the demand for performance grows more slowly than the supply. The demand for high-end PCs will slow as more users are satisfied with their systems.

Memory and microprocessors diverge. The quest for performance led to specialized manufacturing. Microprocessors sought clock speed; memory chips sought only capacity. Now we have fast microprocessors but no fast, big memories that can keep up. The different objectives meant that attempts to improve the performance of the PC met diminishing returns. When the PC was introduced, memory and microprocessor speeds were about the same. Since 1981, memory chips have grown 4,000 times larger, but are only 5 to 7 times faster. Leading-edge microprocessors have gotten 400 to 500 times faster. This widening performance gap between the memory and the microprocessor means that even large improvements in the microprocessor's speed translate to smaller and smaller improvements in the speed you see. Less and less added system performance will slow consumer buying. As the market for high-end PCs declines, critical engineering talent will be diverted from PC improvements to emerging untethered applications.

Embedded microprocessors. More performance per 18-month generation is one side of the coin. The other side is the lower cost per transistor with each generation. High-end applications for leading-edge PCs, ASICs, DSPs, and PLDs paid for process development and for fabs. *This reduced the marginal cost of production for a huge wedge of low-cost, trailing-edge applications*. High-end applications bought and used the then high-end equipment and then moved on to the next-generation equipment. Already-amortized production equipment running a semiconductor process a generation or two behind the leading edge produced low-cost chips. The electric toothbrush, the microwave oven, the hair dryer, and a zillion other applications benefited. The processes and fabs supported by the leading-edge applications enabled, in their wake, the growth of embedded microprocessor and microcontroller shipments from nothing in 1970 to 8.5 billion units in 2000.

Increasing capability and the declining cost of the microprocessor enabled it to displace rivals and to grow its range of applications. The Furby, the blender, and other electronic toys, games, and appliances could not have paid for semiconductor process development or for their own fabs on the strength of individual demand.

The one-million-transistor chip: 1984

PLDs emerge. IBM introduced the 1-Mb memory chip. The million-transistor chip signaled practical utility for the programmable logic device. PLDs began with roles in consolidating logic and in prototyping logic circuits. The PLD's strength is that it is manufactured as a generic chip that is personalized in the field to suit a particular application. The PLD's weakness is its inefficient use of transistors. Abundant transistors magnify the PLD's strength and diminish its weakness.

The microprocessor, the microcontroller, and the digital signal processor have dominated embedded designs for thirty years. They are entrenched in development systems, in legacy applications and in operating systems, in the engineering design community, and in the educational system. They are backed by billion-dollar corporations such as Intel, Motorola, AMD, and Texas Instruments. Similar things might have been said about the mechanical calculator. Times change.

The microprocessor and its kin arose when cost was important, performance had to be adequate, and power consumption was not a consideration. Now power efficiency is important and performance requirements are rising, particularly for battery-powered systems. Now, solution efficiency matters, but the microprocessor's solutions aren't efficient. Imagine a complex equation. IC macros or an ASIC implement the equation directly. A microprocessor implements the equation with a program. The program is a sequence of instructions that simulates the complex equation. Efficiency is lost in translating the equation into the microprocessor's instructions and also in fetching, decoding, and executing the instructions. There's more overhead in the software that manages the microprocessor and its programs. The microprocessor's position is weakening.

The IC macros that the microprocessor displaced won't make a comeback. ASICs won't displace the microprocessor either. The ASIC's efficiency advantage (a direct solution in hardware) is offset by its cost and by its focus on a single application. What's needed is ASIC-like efficiency in a general-purpose component. *For a growing range of applications, that'll be a PLD.*

The production of transistors outpaces our ability to design them. Rising mask costs make diverse chips uneconomical. Both trends encourage building general-purpose chips that can be fitted to an application *after manufacture*. The beneficiaries of this trend are the twin 800-pound gorillas of the PLD business, Altera and Xilinx.

Microcontrollers. Moore's law improvements have led to extreme performance as PC microprocessors or to proliferation as diverse microcontrollers. The microcontroller market requires just-adequate performance at the lowest possible cost. Typical consumer applications, such as electric irons, remote controls, and dishwashers, don't demand more performance or capability than an 8-bit microcontroller supplies. Since they don't need more performance or capability, there's constant pressure to reduce the size, and therefore the cost, of the microcontroller.

Moore's law progress continues to shrink the microcontroller's circuits until the chip's bonding pads (where wires connect the circuit to the outside world) determine its size. When the chip's bonding pads stop the *chip* from shrinking, Moore's law progress shrinks the *circuit* to a small portion of the chip's area. Triscend, Hitachi, and Cypress Microsystems spend transistors in this "free" space to build programmable logic that users configure into custom peripherals. One chip with a microprocessor core and some programmable logic substitutes for a whole range of custom microcontrollers. "Soft" descriptions of processors and peripherals solve the need to redesign for each new process generation. With a combination of soft peripherals and programmable logic, Triscend, Hitachi, and Cypress Microsystems can consolidate the microcontroller market.

The fabless semiconductor model: 1985

Gordie Campbell and others founded Chips and Technologies in 1985. Chips and Technologies made chip sets for PCs, but it didn't own a fab. "Chips" was the first of the fabless semiconductor companies.

Vertical-to-horizontal fragmentation of IDMs. In the days of integrated device manufacturers (IDMs), each company designed and manufactured its own chips from beginning to end. In addition to process development, chip design, and manufacturing, early IDMs made their own wafers and their own semiconductor processing equipment. As the chip industry grew, wafer production and semiconductor processing equipment moved to standards and became independent businesses. Today, the semiconductor industry is fragmenting into intellectual property (IP) designers, intellectual property aggregators, and foundries.

Horizontal fragmentation of an industry raises the efficiency of its engineers. Each IDM once designed its own PCI controller and other common functions. Today, these are readily available for licensing. A few teams now design Ethernet controllers and license them throughout the industry. Horizontal fragmentation of the IDMs enables the small design teams that develop intellectual property "cores" at companies such as Lexra, Inc., inSilicon inc., and Nova Engineering. Horizontal fragmentation also benefits foundries, such as Chartered, TSMC, and UMC.

The Web browser: 1990

By the time the first browser was introduced in 1990, PCs based on Intel's x86 microprocessors dominated the market. The Internet and the World Wide Web connected the world's computers into a universally accessible resource of computing, data storage, access ports, and data transport.

The microprocessor core: 1991

In 1991, ARM Ltd. began licensing the ARM6 microprocessor core. It was the beginning of the transition that is still transforming the industry. Before the ARM6, a microprocessor's value resided in a physical chip. With the ARM6, the value began to move from the physical chip to the *design file that described the microprocessor's design*. The first microprocessor cores were "hard" cores. A hard core describes the layout of the circuit's elements in a particular semiconductor process and is not portable to a different foundry or to a different process.

Hardware gets softer. Instead of designing functions to a particular semiconductor process—called hard cores, designers now build soft cores. A soft core is less efficient than a hard core, but it does not have to be redesigned to move from one process generation to the next (as the hard core does). Hardware is getting softer. First, the value was in the physical chip. Hard cores moved the value to the circuit's description in a particular semiconductor process. Today, soft cores move the value to a portable (across processes) description of the circuit. The softening of hardware benefits providers of soft cores, such as ARM, ARC, Triscend, and Tensilica.

The Palm Pilot: 1996

Tethered and untethered devices. The Palm Pilot signaled a world splitting into tethered and untethered devices. Tethered devices connect to the power grid and constitute a globally interconnected grid of computing, data storage, access ports, and data transport. Untethered devices do not get their power from a wall outlet and must, therefore, make power efficiency a primary goal. Many of these untethered devices (personal digital assistants, GPS receivers, MP3 players, cameras, and cell phones) are consumer items with high-end computing requirements. These requirements make low cost and high performance primary system design goals.

In the early days of computing, humans collected the data, converted it to a format the computer would understand, and delivered it to the computer. Today's untethered computing devices collect their own data at its source. Today's computers are so fast and so numerous that it's not practical for a human to be in the middle of the connection between the data and the computer.

Digital signal processors. The DSP is a microprocessor that is specialized for high-data-rate processing. DSPs and microprocessors suffer from a common problem in untethered devices. Historically, DSPs and microprocessors increased performance by increasing the clock rate. Generally speaking, doubling the clock rate doubled the performance; but it also doubled the power use. That's a problem for untethered devices that have high computing requirements and also must conserve power.

DSPs are the fastest-growing segment of the microprocessor market today, but they are doomed. Today's untethered devices contain a DSP and a microprocessor. Two processors is one too many; the microprocessor is the housekeeper for the system, so it will stay. The DSP's functions will migrate to the microprocessor.

Untethered devices move to x86. The end of the DSP in untethered devices isn't the end of the story. ARM Holdings, Ltd. dominates microprocessors for untethered devices with about 70% of the market. Its position seems secure. Is it? I don't think so. Everything is connected to the global information grid. The content of the global information grid was built by the x86-based PC. The simplest way to achieve file, browser, and application compatibility with the global information grid is to provide an x86 in the untethered device. Then we're back to two microprocessors in an untethered device. One will have to go. Will ARM or x86 dominate untethered devices?

An enormous repository of legacy software (protocol stacks, operating systems, and applications) ties unteth-

What Has Happened in Semiconductors
Trends and Consequences

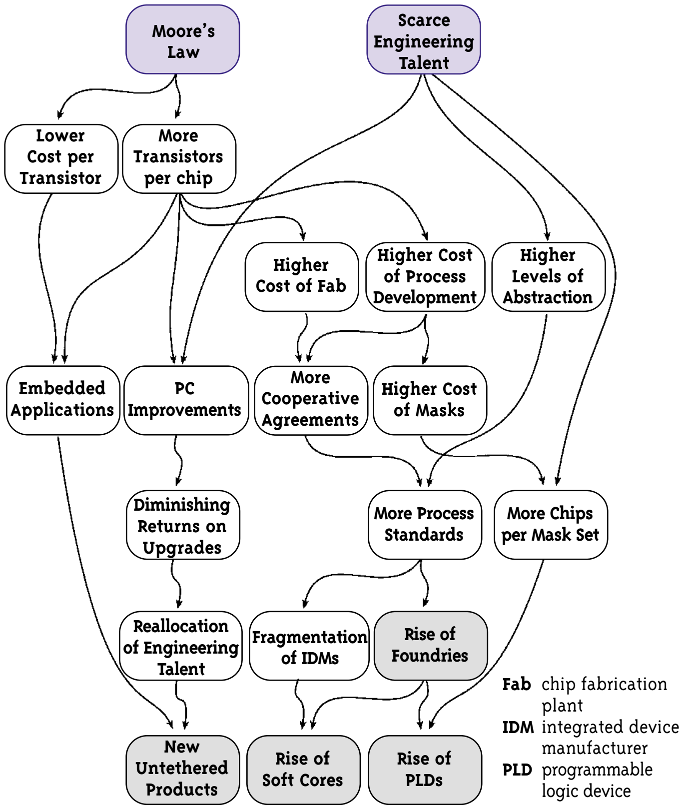

ered devices to the ARM microprocessor. A similar repository of legacy software ties the global information grid to the x86 microprocessor. But there's a huge difference that will determine the outcome. The global information grid is interconnected; untethered devices are isolated. Untethered devices connect to the global information grid and they connect to each other through the grid. Untethered devices can migrate from ARM to x86, but the global information grid cannot move from the x86.

Lessons

Moore's law and the scarcity of engineering talent drive the semiconductor industry. As they do this, four themes recur: changes in supply and demand, proliferation and consolidation, vertical-to-horizontal fragmentation, and rising levels of abstraction. Moore's law improvements, in lowering the cost per transistor and in putting more transistors on a chip, increase the number of potential applications. Moore's law improvements require more complex processing, raising the cost of fabs and of process development. The higher cost of fabrication plants and of process development leads to cooperation among chip makers. Cooperation and the higher levels of abstraction that conserve scarce engineering talent encourage the rise of semiconductor process standards. Process standards foster horizontal layering of the industry and promote foundries. Foundries contribute to the rise of soft cores and to the rise of programmable logic. The future is new untethered products built of soft cores on high-volume programmable logic devices.

Brion N. Shimamoto

Nick Tredennick and Brion Shimamoto
April 11, 2002

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Altera, Xilinx	Fabless	Excellent	Builds PLDs as the industry moves to PLDs. Dominant share of a rapidly growing market. Uses soft cores as the industry moves to soft cores. Fabless as the industry moves to fabless.
ARM Holdings, Ltd.	Fabless	Excellent	Dominates the market for soft cores as the industry moves to soft cores. Fabless as the industry moves to fabless.
TSMC, UMC	Foundry	Excellent	Dominant foundry in an industry moving to foundries.
ARC Cores, Tensilica	Fabless, Startup	Good	Builds soft cores as the industry moves to soft cores. Fabless as the industry moves to fabless. Builds energy-efficient configurable microprocessors as the industry moves to power conservation.
Celoxica	Software, Startup	Good	Raises the level of abstraction for designers. Increasing the pool of designers. Builds soft cores for PLDs as the industry moves to soft cores and to PLDs. Must overcome entrenched logic design biases.
Chartered	Foundry	Good	Third largest foundry in an industry moving to foundries.
Cypress Microsystems, Triscend	Fabless	Good	Exploits programmable logic to consolidate the microcontroller business. Microcontrollers are a low-margin business. Early in developing market.
Actel, Lattice, QuickLogic	Fabless	OK	Builds PLDs as the industry moves to PLDs. Minority share of the rapidly growing PLD market. Uses soft cores as the industry moves to soft cores. Fabless as the industry moves to fabless.
Applied Materials, Credence Systems, Lam Research	Systems	OK	Equipment makers supply the foundries and they supply the integrated device manufacturers. Business is assured but subject to the industry's boom and bust cycles.
inSilicon, Lexra, Nova Engineering	Fabless	OK	Builds soft cores as the industry moves to soft cores. Early in a developing market.
Transmeta, VIA Technologies	Fabless	OK	Fabless as the industry moves to fabless. Builds power-efficient microprocessors as the industry moves to power conservation. Builds PC microprocessors as buyers move to lower-margin systems.
AMD, Intel	Integrated	Struggle	Builds high-performance microprocessors as the industry moves to power-efficient microprocessors. Integrated device manufacturer as the industry moves to soft cores and to foundries.
Compaq, Dell, Gateway	Systems	Struggle	PC buyers will move to lower-margin systems as system performance produces diminishing improvements on levels that already exceed needs of large market segments.
Fujitsu, LSI Logic, NEC, Toshiba	Integrated	Struggle	ASIC suppliers lose business to PLD companies as the industry moves to more generic chips. Integrated device manufacturer as the industry fragments horizontally. Libraries are hard cores as the industry moves to soft cores.
Motorola, Texas Instruments	Integrated	Struggle	Builds digital signal processors as the industry moves to microprocessors with vector extensions and to PLDs. Integrated device manufacturer as the industry moves to soft cores and to foundries.

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	4/8/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	22.40	14.66 - 33.60	8.7B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	42.78	29.00 - 53.30	15.6B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.60	4.76 - 6.38	£86M
ARM Limited (ARMHY***)	Microprocessor and System-On-A-Chip Cores	11/26/01	16.59	12.08	8.39 - 19.20	4.1B
Calient (none*)	Photonic Switches	3/31/01				
Celoxica (none*)	DKI Development Suite	5/31/01				
Cepheid, Inc. (CPHD)	MEMS and Microfluidic Technology	12/17/01	4.73	3.49	1.48 - 11.48	93.6M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	26.05	16.06 - 34.0	3.6B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	22.50	14.00 - 28.95	2.7B
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
QuickSilver Technology, Inc. (none*)	Dynamic Logic for Mobile Devices	12/29/00				
SiRF (none*)	Silicon for Wireless RF, GPS	12/29/00				
Taiwan Semiconductor (TSM†)	CMOS Semiconductor Foundry	5/31/01	14.18††	19.58	8.39 - 20.99	65.9B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	2.91	1.17 - 25.00	395.2M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	10.20	4.25 - 10.71	27B
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	13.30	9.71 - 29.25	1.0B
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	40.64	19.52 - 52.14	13.7B

† Also listed on the Taiwan Stock Exchange

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

* Pre-IPO startup companies.

** ARK is currently traded on the London Stock Exchange

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

NOTE: This list of Dynamic Silicon companies is not a model portfolio. It is a list of technologies in the Dynamic Silicon paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. The authors and other Gilder Publishing, LLC staff may hold positions in some or all of the companies listed or discussed in the issue.

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

www.dynamicsilicon.com