Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

The Investor's Guide to Breakthrough Micro Devices

The Tech Sector Rocks

Technology is the best sector in the market. Moore's law predicts sustained improvements in integrated circuits (ICs). ICs get cheaper *and* they get more capable. As they get cheaper, they invade new areas. As they get more capable, they invade new areas. Carpets, automobiles, fabrics, tires, soft drinks, and entertainment don't do this. Markets for these things grow with improvements in productivity and they grow with the invention of new materials. But, tires aren't going to invade the carpet business or the soft drink business. Marketing can drive soft drinks into more refrigerators or into more countries, but there's no Moore's law equivalent that's going to drive them into automobile engines or into light bulbs.

Some investors shun the tech sector for its boom and bust cycles (May 2001, *Dynamic Silicon*). This year may be one of the worst ever in the semiconductor industry. Following on two years of better than 30% growth, the industry may *shrink* by 20%. There's nothing new about this cycle; it's old hat for industry veterans. The cycle is a function of the way the semiconductor industry does business. In spite of these cycles, the semiconductor industry has sustained a cumulative growth rate of 17% for forty years. Farm equipment can't do that; building materials can't do that. Moore's law works for integrated circuits; it does not work for tractors. A major plus for semiconductors: there's little regulatory interference and little political blundering. The semiconductor industry's fortunes aren't tightly tied to the FCC, FDA, FTC, PUC, or other three-letter agencies. When interference and blundering cripple the telecommunications business, semiconductors suffer—but the effects are secondary.

The line below the banner on this page, reads, "The Investor's Guide to Breakthrough Micro Devices." My job and this line strike me as ironic. I'm not wealthy, I don't have a financial background, and I'm no day-trader (I probably average a trade a year). What am *I* doing writing an "investor's guide?" First, the simple explanation: you and I are technology enthusiasts; you are a technology investor; I am a technology guide. I have a technical background and decades of technical experience. You would like Nasdaq symbols for companies with breakthrough ideas, sustainable advantages, growing markets, and great management. You would like to see high growth in the stock prices of listed companies. Against that model, I'm recommending companies such as Calient, Celoxica, QuickSilver, SiRF, Triscend, and Tensilica that aren't even public. What kind of advice is that when your guide is pointing at companies you cannot invest in?

The industry downturn doesn't mean that ICs have saturated. It's not as if we've sold a car to every driver and have begun to live on replacements. That won't happen in the next ten years. It may not happen in the next twenty years. Moore's law ensures a replacement market in ICs as cheaper, better chips displace old ones. And ICs will continue to find new markets. I may have to change the title to *Dynamic Carbon* (the element base for nanotechnology), and I may have to talk about nanotechnology, but it will be technology. I'm not going to be talking about pork bellies. In good times, novelty and performance drive electronic content. In bad times, cost and efficiency drive electronic content. In good and in bad times, semiconductors continue their invasion.

Write the alphabet on a sheet of paper; find a semiconductor stock for each letter; you have a long-term investment portfolio that will grow at 17%. Learn something about technology and you can do better. You can separate moribund companies from leading-edge companies. Here's my reasoning for the Dynamic Silicon Companies.

In This Issue:

Vol. 1, No. 8
August 2001

ARC Cores, Tensilica, Transmeta

Microprocessors are ubiquitous. The microprocessor invaded embedded systems (electronic components you don't realize are there) in 1969. It invaded the computer in 1974. Since its commercial introduction by Intel in 1971, microprocessor shipments have grown to billions of units per year. About 150 million microprocessors are the central processor in a computer system; the rest are in embedded systems (e.g., blenders, hair dryers, rice cookers, cell phones, and cameras). The schools, the engineering community, and countless installed development systems back microprocessor-based designs. The microprocessor won't be going away soon. Moore's law and the microprocessor's success create opportunities.

Since its introduction, the microprocessor has been an IC with a fixed instruction set. Each microprocessor manufacturer designs this instruction set for a range of applications. Engineers buy the microprocessor with the instruction set that best fits their application. It's always a compromise.

A stock microprocessor has most of what's needed, but lacks features key to, say, an automatic transmission's gear-shifting algorithms. ARC Cores (ARK) and Tensilica are Dynamic Silicon Companies because they offer relief to the engineer shoehorning the automatic transmission's algorithms into a stock instruction set. ARC and Tensilica offer configurable soft cores rather than hard ICs. A soft core is independent of the semiconductor process used to make the chip—allowing the designer to choose a chip manufacturer. ARC and Tensilica offer more than soft cores; the engineer configures the microprocessor and can augment the instruction set. Using design tools from ARC or Tensilica, the engineer tailors the microprocessor's performance to meet requirements. No sense building a multiply-accumulate unit into the chip if the automatic transmission or hair dryer doesn't need one. In addition, the engineer adds custom instructions to suit a particular application. A few application-specific instructions can return huge benefits in performance or in energy savings.

The PC grew up with Intel's x86 instruction set. Most personal computers run a version of Microsoft Windows plus a word processor and a spreadsheet. The x86 instruction set is not a perfect fit for the PC's applications, but we're stuck with it. Transmeta found a way to get around the PC's restrictions. Transmeta offers microprocessors with two instruction sets. From the *outside*, Transmeta microprocessors have an x86 instruction set. Under the hood, a layer of software hides the real instruction set. This layer between the real hardware and the outside world monitors the microprocessor's behavior and dynamically adjusts the microprocessor's operation to suit conditions. A Transmeta microprocessor might notice that it's not working hard and lower the clock frequency and the power-supply voltage to reduce the energy drain. Since there's a software layer between the outside world and the hardware, a microprocessor in the field could learn new instructions—a hardwired microprocessor, such as those from Intel and AMD, couldn't. A Transmeta microprocessor might learn x86 code "signatures" and use Transmeta shortcuts for lengthy operations. Transmeta is a Dynamic Silicon Company for the power of the software layer between its hardware and its external x86 instruction set.

Photo courtesy of Triscend Corporation

Triscend, Cypress

The first microprocessors had to make every transistor count. Moore's law doubled the number of transistors on the chip every eighteen months. Soon, peripheral chips (e.g., ROM, RAM, serial ports) moved onto the chip with the microprocessor, making the chip a "microcontroller." Microcontrollers reduced the parts count in embedded applications. But microcontroller types multiplied into the thousands, reflecting the diversity of embedded applications. Because there are so many microcontroller variations, engineers have difficulty finding the right microcontroller for a particular application.

Triscend offers an alternative that exploits Moore's law increases in on-chip transistors. Triscend offers microcontrollers with programmable logic on the chip. An engineer at a PC drags and drops peripheral functions from a menu to create a *custom* microcontroller. A few standard chips, together with development software and with a library of

DynamicSilicon

Editors	**Nick Tredennick**
	Brion Shimamoto
Publisher	**Lauryn Franzoni**
Web Editor	**Jorin Hawley**
Designer	**Julie Ward**
President	**Mark T. Ziebarth**
Chairman	**George Gilder**

DynamicSilicon 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 Single-issue price: $50. For subscription information, call 800.229.2573, e-mail us at dynamicsilicon@gilder.com, or visit our website at www.dynamicsilicon.com

soft IP (intellectual property) peripheral functions, displace thousands of custom microcontrollers. Cypress Microsystems, a wholly owned subsidiary of Cypress Semiconductor (CY), also builds microcontrollers with configurable peripheral functions. These Dynamic Silicon Companies take advantage of Moore's law and of the softening of hardware. Triscend is fabless and Cypress owns fabs. They provide customized microcontrollers without fragmenting IC manufacturing.

TSMC, UMC, Chartered

A semiconductor foundry costs $2 to $3 billion to build and the equipment is quickly dated. With huge fixed costs, it's best to keep the foundry running at 100% capacity to amortize fixed costs over the highest production volumes. Average capacity is expected to be at 50% to 60% this year, so the foundries are losing money. That's the short-term view.

The long-term question is whether the foundry model is correct. Traditionally, large companies like Intel, IBM, Hitachi, Motorola, TI, and STMicroelectronics designed and manufactured their own chips. These are the integrated device manufacturers (IDMs). IDMs thought of the foundries as second-tier suppliers to design houses too poor to afford their own "fab" (semiconductor fabrication line). Foundries ran a generation or two behind the IDM fabs. If the IDMs were at 1.0-micron line widths, the foundries were running at 1.5 or at 1.2 microns. Not so today. Taiwan's world-leading foundries, TSMC (TSM) and UMC (UMC), have world-leading semiconductor processes as well. If the semiconductor processes are equivalent, then the answer to which model is better depends on other things.

Intel has plants dedicated to Pentium microprocessors. There's no stopping to change the production line for a different product. That's the ideal way to run a fab. If, however, demand for Pentiums wanes, the fab runs at less than capacity. At an independent foundry, demand for cell phones might be up when demand for personal computers is down. Capacity idled by a downturn in automotive chips might be filled by growth in entertainment electronics. Filling a foundry with odds and ends has disadvantages too: the production line has to be set up differently for different products. It takes time and it costs money to reconfigure the fab's production line. Sometimes customer variety is an advantage, but when the whole industry is down, no one runs at capacity.

The issue is over whether a vertical (IDM model) or a horizontal (foundry model) organization is the industry's future. "Foundries" separate the designers from the manufacturers. The foundry model succeeded because it enabled a new class of design companies. Companies too small to build a fab could now design chips and have them built at an independent foundry. As the value in chips moves from circuit hardware to circuit descriptions (soft IP), the foundry model becomes superior to the IDM model. *The chip industry is splitting into IP developers, IP integrators, and foundries.* ARC Cores, Celoxica, and Tensilica are IP developers. With a horizontally organized industry, almost any company can develop and license IP; almost any company can be an IP integrator. TSMC, UMC, and Chartered are the leading foundries.

In this model, Altera (ALTR) and Xilinx (XLNX) are meta-IP developers. They develop "hard" IP that is the basis for hardware chips manufactured at TSMC and at UMC. Intel, with an IDM model, builds hard microprocessor chips—think of Intel as making millions of perfect copies of a painting that has intrinsic value. By contrast, Altera and Xilinx, with a foundry model, build hard programmable logic chips—think of them as millions of white boards. Value accrues to the programmable logic chip (white board) when it is "personalized" (painted) by loading the soft IP.

Altera, Xilinx

Because of its success in the x86 microprocessor market, Intel has been one of the most profitable companies of all time. What's the next Intel? My answer: Altera and Xilinx. Both? Won't one wipe out the other? How did you select that market segment? "Both" *is* the right answer. One won't wipe out the other. And programmable logic is the right market segment. Best news first! Altera and Xilinx have been at each other's throats almost since their businesses were founded. They're two mice fighting for control of a cheese factory, struggling to own a market that's too big for both of them. In July 2001, they signed a royalty-free patent cross-license agreement and a peace treaty. No more patent litigation for five years! Five years seems too short, but it's a start. As someone who has watched for years as these companies wasted millions of dollars a month on lawyers and on legal proceedings, I am euphoric. The engineers, executives, consultants, and retirees who have been mired in this mess for years can now do something productive.

Conference agendas pit ASIC (application-specific integrated circuit) advocates against PLD (programmable logic device) advocates; vendors debate cost, circuit capacity, performance, and time-to-market. It's entertaining, but the points are valid only for the performance corner the ASIC vendors are backing into. The real battle is over the level of abstraction. It's been fought before. Circumstances haven't changed, so the result will be the same. The battle was first fought between assembly language programming and high-level language programming. The conference debates were similar: performance and code size. The result is clear today: no one can afford assembly language programming. A high-level language raises the level of abstraction and thereby increases the pool of qualified people and improves their productivity. For ASICs, circuits are fixed (static) in the hardware; build the ASIC and it works or it doesn't work. SRAM PLDs have logic circuits and wires on one chip-layer and have static memory (SRAM) on another layer. Bits in the SRAM make connections between the logic circuits and wires to build circuits. For SRAM PLDs, circuits are bit patterns; download the "personalization file" into the PLD and see if it works. If it does, great; if it doesn't, modify the file and download it again. Even systems in the field can be modified this way. With the right codes and an Internet connection, engineers can modify systems in the field from their offices—no truck-rolls to mess with the hardware.

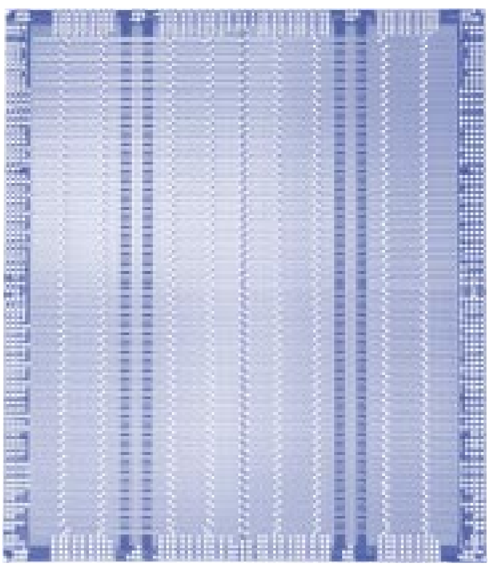

Photo courtesy of Altera Corporation

Fig. 1. An Altera APEX EP20K1500 chip. The tens of millions of transistors are so small that only grand-scale layout features are visible. Dots in vertical rows and around the edges are "flip-chip" pads.

Altera and Xilinx have built their SRAM PLD businesses on *displacing ASICs* and on *prototyping*. The average ASIC design gets a little larger each year. Moore's law is growing PLDs much faster. That means PLDs are cheap enough and fast enough to capture more of the ASIC market each year. Circuit prototyping continues to be a large and growing market for PLDs. Altera and Xilinx have grown their respective businesses at about 35% per year for at least the last ten years. That's for operating income, net income, assets, and even R&D expenses. They've grown at this rate on a strategy of ASIC displacement and of prototyping. That's about as fast as it's reasonable to grow a business and have it stay healthy. If your company grows 35% per year, almost half of your employees have been there less than two years. You might grow a meatpacking business that fast, but this is an engineering business developing leading-edge silicon and software. Development requires skilled engineering teams working together for years, who are supported by extensive computing resources. SRAM PLD makers haven't been able to grow fast enough to test the elasticity of the market, so margins are high.

Altera and Xilinx are big now. They're looking for other markets. New markets are easy to find: the *microprocessor* and *DSP* (digital signal processor) markets together are about ten times the size of today's PLD market and they are ripe. Funneling everything through a microprocessor or DSP instruction set isn't efficient. BlueArc demonstrates the superiority of SRAM PLDs over microprocessor-based processing in file servers. Chameleon Systems demonstrates the ability of SRAM PLDs to implement reconfigurable base stations (for cell phone networks). IP Semiconductors designed a network processor, called SpeedRouter, as a soft IP core that fits in a third of a Xilinx Virtex SRAM PLD. The SpeedRouter is fast enough to process data packets without storing them between processing stages and its size scales to meet requirements. As Xilinx shows in recent advertising, SRAM PLDs can out-process DSPs in their home application of signal processing. SRAM PLDs are beginning to displace microprocessors and DSPs in *performance-oriented* systems. Once the engineering design community gets used to the idea that it's feasible and practical, these designs will proliferate. Migration from implementations based on microprocessors and DSPs to implementations based on SRAM PLDs will be slowed by the need to reeducate an engineering design community steeped in thirty years of microprocessor-based design.

In the x86 microprocessor market, Intel hasn't had a strong competitor. Since it dominated the market, there was little real incentive for Intel to improve the microprocessor's performance rapidly. Fortunately for consumers, Intel *thought* it heard the footsteps of the workstation vendors with their RISC microprocessors, which caused Intel to push the development of its x86 microprocessors. Altera and Xilinx dominate the PLD market with a little more than a third each, leaving less than a third to be divided among others. The competition between Altera and Xilinx benefits consumers: PLDs and their development systems and support will improve faster than they would in an environment dominated by a single company. It's better for Altera, better for Xilinx, and better for the industry as the SRAM PLD makers compete to open new markets.

Celoxica, Wind River Systems

In custom logic, such as in an ASIC, the circuits and the algorithms are “hard.” That is, they are built into the physical structure of the chip and cannot be changed. An ASIC, therefore, serves a single application. The value is in the component. In a microprocessor, the circuits (including the instruction set) are hard and the algorithms are “soft.” The algorithms are soft because they exist as programs that can change (hence “software”). A microprocessor serves a range of applications. The intrinsic value resides both in the microprocessor and in the programs. In an SRAM PLD, both the circuits and the algorithms are soft. The value is in the description of the algorithm and of the circuit, both of which are soft.

The value in a design was once its hardware; today a design’s value is in its soft description—its intellectual property. Celoxica’s DK1 development tools and Handel-C design language help engineers map algorithms into IP for SRAM PLDs. Moving design from hardware description languages to a variant of the C programming language raises the level of abstraction and it greatly expands the pool of qualified designers. Engineers skilled in C are more common than engineers skilled in hardware description languages. Celoxica also provides design services. Celoxica is a Dynamic Silicon Company for its strategy to be a supplier of IP.

Celoxica, Wind River Systems (WIND), and Xilinx are defining interfaces that enable engineers to move functions into SRAM PLD-based accelerators in embedded systems. Wind River’s VxWorks is the leading operating system for embedded systems. Wind River complements its Tornado development environment with Celoxica’s DK1 development tools, to enable the development of reconfigurable systems. Xilinx, Wind River, and Celoxica are building a “bridge” between the microprocessor and the SRAM PLDs on a circuit board. This bridge permits compute-bound functions to migrate seamlessly from the microprocessor to a more direct implementation inside an SRAM PLD. Instead of defining new interfaces, Wind River should consider the resource manager from QuickFlex, a pre-IPO startup. QuickFlex’s resource manager complements Wind River’s VxWorks operating system in the dynamic allocation (“bridging”) of logic to the board’s SRAM PLDs.

Analog Devices, SiRF, Coventor

The world is splitting into tethered and untethered devices. Tethered devices use wall sockets for power and they build the global information grid. The grid is the world’s connected repository of computing, access ports, data transport, and storage. Untethered devices connect to the grid by radio and they do not need wall sockets. Mobile devices, a subset of untethered devices, travel with us as the collectors and consumers of data. Mobile devices translate the sights, sounds, smells, and actions of the real world into the ones and zeroes of the digital information world. Mobile devices, such as the cell phone, the personal digital assistant, and the GPS receiver, are in the overlap of the zero-cost, zero-power, and zero-delay segments of the electronics market (January 2001, *Dynamic Silicon*). I call this segment the leading-edge wedge because of its conflicting requirements for low cost, long battery life, and high computational capability.

The cell phone is the most challenging example. Market pressures work to reduce its cost, to lengthen talk and standby times, and to improve its performance. Today’s cell phone, which is primarily a radio with RF, IF, and baseband sections, requires at least a hundred components (July 2001, *Dynamic Silicon*). Reducing the component count in the cell phone is particularly challenging because many of today’s discrete components (e.g., microphones, antennas, inductors, switches, relays, tuning capacitors, and oscillator components) do not flatten well for placement onto a chip. Analog Devices (ADI) is a Dynamic Silicon Company. Analog Devices makes analog RF and IF components and chip sets, direct-conversion receivers (which eliminate the receiver’s IF section), digital signal processors, data converters, and even MEMS (microelectromechanical systems) accelerometers. In short, it is positioned to supply all critical components to the rapidly growing mobile-device market.

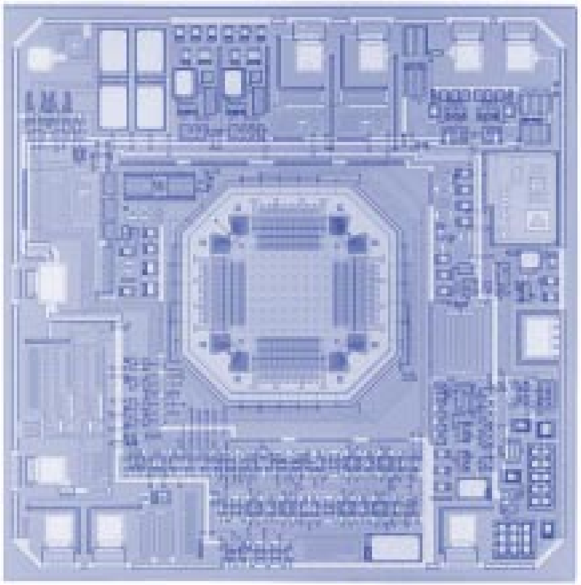

Photo courtesy of Analog Devices, Inc.

Fig. 2. Die photo of a two-axis accelerometer from Analog Devices. This ADXL202 is a low-cost, high-volume MEMS chip.

Mobile devices will know the time and will keep track of where they are. SiRF is a Dynamic Silicon Company because it offers consumer-focused GPS receivers both as chips and as soft IP. SiRF's designs will be in everything that moves—from wristwatches to tractors.

Photo courtesy of SiRF Technologies

Fig. 3. SiRF's GPS-receiver reference design.

Cahners In-Stat Group (www.instat.com) estimates that the market for RF MEMS will grow from $1 M in 2001 to $350 M by 2006. With better than 220% CAGR, where are the RF MEMS companies that will take advantage of this opportunity? Many of them are large companies, such as TI, Analog Devices, Intel, Motorola, and STMicroelectronics, that supply cell-phone components. The company these suppliers turn to for RF MEMS expertise is Coventor. Coventor supplies IP, professional services, development tools, and connections to MEMS foundries. Coventor is a Dynamic Silicon Company for its key position between the cell-phone component suppliers and the MEMS foundries. Coventor's strategic focus is MEMS for wireless, for optical, and for life sciences applications—the three highest growth segments of the MEMS market.

QuickSilver

The cell phone is the highest-volume mobile device. It is the quintessential leading-edge-wedge device. Its market seeks longer talk times, access to email and other services, and location-accuracy of fifty feet. The cell phone's baseband section contains, among other components, ASICs, a microprocessor, and a DSP. As the cell phone moves from second- to third-generation services and protocols, computing requirements increase at greater than a Moore's law rate, *forcing more functions into custom hardware*. The too-general-purpose microprocessor and DSP struggle to keep up; they remain in the designs, but they do less and less of the work.

One component of microprocessor and DSP performance is clock frequency. Since the introduction of the PC in 1981, with its 4.77-MHz microprocessor, microprocessor clock frequencies have increased at 35% per year. Today's fastest microprocessors run at 1.8 GHz. This creates a dilemma for the microprocessor and for the DSP because power dissipation increases directly with clock frequency: more performance implies more energy consumed. The market for mobile devices wants higher performance and *less* power.

The microprocessor and the DSP will evolve from hard ICs to soft IP that can be integrated with other functions to reduce the chip count overall. That transition will help, since it will reduce signal distances, wiring, chip packages, and driver (signal amplifier) sizes. It won't be enough; the problem is deeper—software implementation of functions is too inefficient. ASICs implement direct solutions, so they are efficient. But ASICs are too expensive. ASICs also take too long, and they lack flexibility. Programmable logic solutions can implement direct solutions and they are flexible, but they have high overhead and they are expensive.

The above dilemma is QuickSilver's opportunity. QuickSilver's Adaptive Computing Machine (ACM) combines the speed of the ASIC's direct implementation with the flexibility of programmable logic. Because QuickSilver's chip is built for specific applications, however, it doesn't incur the overhead of general-purpose programmable logic components. QuickSilver offers the ultimate in softening of hardware for mobile applications. Functions are "paged" into hardware resources as they are needed. QuickSilver's unfortunate choice of the ACM name confuses the issue by using the software-oriented word "computing" to describe a *paged-hardware* implementation. Direct hardware implementation (via an ASIC, for example) offers the ultimate speed alternative to the "computed" solutions offered by executing instructions on a microprocessor or on a DSP. QuickSilver's ACM, for example, is superior to both: its resource manager pages in custom (soft) hardware for each task. QuickSilver's ACM can adapt to cell-phone protocol changes in the field.

QuickSilver faced three major challenges: acceptance in the venture community, weaning one of the few major manufacturers from traditional solutions to a new idea, and collecting an engineering team that could implement a dynamic-logic chipset. QuickSilver struggled for funding, but finally encountered Techfarm's visionary, Gordon Campbell, whose company Chips and Technologies pioneered the foundry ("fabless semiconductor") model. "Gordy" has become an enthusiastic proponent of dynamic logic. Given the pedigree of its backers (including Techfarm, BellSouth, and Kyocera) QuickSilver has likely made headway with at least one major manufacturer. The engineering team is a more difficult call. The vision at the top is clear and consistent, but implementation is details and compromises. Implementing the pure vision would surely lead to the "Illiac syndrome," which simultaneously pushes innovations on too many fronts and is thereby doomed to failure. Instead, QuickSilver's chipsets are likely to walk a few innovations at a time, through several generations, to the ultimate vision.

Calient

The network core will get dumber. There will be too much traffic for it to be anything else. Converting photons to electrons, sorting packets, and converting the electrons to photons again in an optical-electrical-optical (OEO) switch creates a bottleneck because electrons are slower than photons. The "E" in OEO has to "see" what's in the light. That takes time. Instead of converting and sorting packets at each node, traffic will be sorted and "groomed" near the edges of the network where the information flow is still slow enough for electronic processing. Groomed and aggregated bundles of signals will route through Calient switches on optical interstate highways. Calient's DiamondWave router bounces the signal from each fiber on one port off MEMS mirrors that move to aim the light at a destination fiber. Calient's switch is all optical, called OOO; it doesn't convert incoming signals to electrons. An OOO switch makes routing independent of how (format, data encoding, modulation) information is carried by the light. The switch doesn't care how many lambdas are in the light from the fiber and it doesn't care how the lambdas carry their information. That's important because, unlike operations in an OEO switch, data rates and wavelengths can change without changing the OOO switch.

Where are the MEMS companies?

Analog Devices, Cypress, and Calient make and sell MEMS, but they do other things that dwarf their MEMS business. Where are the MEMS companies that are the pure plays in MEMS corresponding to Altera's pure play in PLDs? There aren't any for the list. First, the MEMS business is a garage-shop operation with every company going its own way. There are no standards for how to make and use MEMS. Each application is unique. Until more high-volume components emerge, there won't be focused component producers. Expect this to take several years. Second, MEMS tend to be enablers rather than end products. MEMS switches, resonators, inductors, and tunable capacitors will enable better cell phones with fewer components, but cell phone chip makers will work with companies like Coventor to add MEMS to their chipsets.

The tech sector

Horizontal specialization will displace integrated device manufacturers as hardware softens. The value in intellectual property is migrating from the integrated circuit chip to the intellectual property added to a standard chip as a final processing step. The integrated device manufacturers, which *tie* circuits to chips, are like the music industry, which welds content and media. MP3, a digital format for music, will break the weld between the media and the content in spite of the music industry's struggle to prevent it. In much the same way, the IC industry is splintering into IP makers, IP integrators, and foundries. It is pre-IPO companies like Celoxica, Coventor, QuickSilver, SiRF, Tensilica, and Triscend that will lead the industry to new business models and not the large companies that have investments in integrated device manufacturing. These are the companies of the future. Even though you cannot invest in these pre-IPO companies today, they show us the model for the industry's future.

The tech sector rocks because semiconductors will invade everything—good times or bad. The semiconductor industry's intrinsic boom and bust cycles make it either a darling or a goat to the press and to short-term investors, but there's nothing better for the long term.

Nick Tredennick and Brion Shimamoto
August 16, 2001

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, several companies on this list may be startups. We will have much to say about these companies in future issues.

Chartered Semiconductor (CHRT www.csminc.com)

Chartered is the third largest chip foundry behind TSMC and UMC. It has a leading-edge semiconductor process and is well positioned to benefit as hardward softens.

Wind River Systems (WIND www.wrs.com)

Wind River is the leading real-time operating systems company for embedded systems. It is positioning itself to benefit as dynamic logic invades embedded systems.

Coventor (pre-IPO www.coventor.com)

Coventor provides development systems, expertise, and IP in the rapidly growing areas of optical, biomedical/chemical, and RF MEMS.

Technology Leadership	Company (Symbol)	Reference Date	Reference Price	7/31/01 Price	52-Week Range	Market Cap.
General Programmable Logic Devices (PLDs)	Altera (ALTR)	12/29/00	26.31	30.06	18.81 - 67.13	11.7B
Dynamic Logic for Mobile Devices	QuickSilver Technology, Inc. (none*)	12/29/00				
MEMS Foundry, Dynamic Logic	Cypress (CY)	12/29/00	19.69	27.28	13.72 - 49.94	3.1B
RF Analog Devices, MEMS, DSPs	Analog Devices (ADI)	12/29/00	51.19	46.00	30.50 - 103.00	17.7B
Configurable Microprocessors	ARC Cores (ARK**)	12/29/00	£3.34	£0.98	£0.48 - 4.29	£499M
General Programmable Logic Devices (PLDs)	Xilinx (XLNX)	2/28/01	38.88	40.00	29.79 - 92.50	13.1B
Configurable Microcontrollers (Peripherals)	Triscend (none*)	2/28/01				
Silicon for Wireless RF, GPS	SiRF (none*)	12/29/00				
Microprocessor Instruction Sets	Transmeta (TMTA)	12/29/00	23.50	2.41	2.37 - 50.88	391M
Photonic Switches	Calient (none*)	3/31/01				
DKI Development Suite	Celoxica (none*)	5/31/01				
Design Environment Licensing for Configurable Soft Core Processors	Tensilica (none*)	5/31/01				
CMOS Semiconductor Foundry	Taiwan Semiconductor (TSM†)	5/31/01	19.86	16.30	11.52- 26.79	18.1B
CMOS Semiconductor Foundry	United Microelectronics (UMC†)	5/31/01	10.16	8.44	6.14 - 13.21	18.1B
CMOS Semiconductor Foundry	Chartered Semiconductor (CHRT)	7/31/01	26.55	26.55	21.05 - 87.75	4.0B
Embedded Operating Systems	Wind River Systems (WIND)	7/31/01	14.32	14.32	12.95 - 50.63	1.2B
MEMS IP and Development Systems	Coventor (none*)	7/31/01				

* Pre-IPO startup companies. ** ARK is currently traded on the London Stock Exchange † Also listed on the Taiwan Stock Exchange

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.