Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

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

Microprocessors In Transition

In the beginning, there was the lowly embedded microprocessor, whence came the microcontroller, the CPU microprocessor, and finally the overrated digital signal processor. Embedded microprocessors (embedded as in a fax machine) and microcontrollers (an embedded microprocessor plus other stuff on the same chip) are designed for low cost. CPU microprocessors and digital signal processors are designed for performance. DARPA (David Tennenhouse, *Communications of the ACM,* May 2000), puts the number of embedded microprocessors at 281 million units. Microcontrollers number 7,257 million units, CPU microprocessors number 150 million units, and digital signal processors number 600 million units. Of course, markets with hundreds of millions or billions of units can't have just one focus for business, but I'll treat these markets as if that is the case, so that we can see where business has been good and where the business will go in the future. I'll identify the large-scale trends that drive the microprocessor industry (there can be profitable niches that buck these trends).

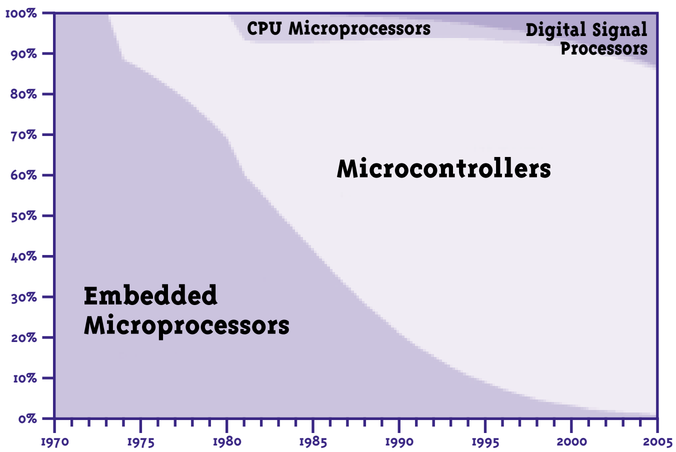

Fig. 1. Microcontrollers dominate microprocessor unit shipments. CPU microprocessors are a small share of unit volumes, but represent about half of the industry's annual revenue.

Fig. 1 shows the percent of unit volumes for *five* categories of microprocessors. Market segments, by year of commercial introduction, are: embedded microprocessors (1971), microcontrollers (1974), CPU microprocessors (1981), workstation microprocessors (1982), and digital signal processors (1983). (Workstation microprocessors, such as Sun's SPARC, are plotted between CPU microprocessors and digital signal processors in fig. 1, but their unit shipments are invisible on this scale.) Since embedded microprocessors, microcontrollers, and digital signal processors are designed for embedded applications, embedded applications have dominated the microprocessor market in unit volumes. Embedded microprocessors began as 100% of the microprocessor market and have shrunk to just 2%, as they have been displaced by microcontrollers with more peripheral functions and by digital signal processors with better computing capabilities.

Embedded microprocessors

The leading embedded microprocessors have been off-the-shelf, general-purpose chips from integrated device manufacturers such as Motorola, Intel, IBM, STMicroelectronics, Hitachi, and Fujitsu. Integrated device manufacturers have instruction-set designers, engineering teams for logic design and circuit layout, manufacturing to

In This Issue:

This month we examine changes in the microprocessor industry—who made money before and who makes money today. Comprehending the industry helps us see its future. The microprocessor business is changing from a vertically integrated business to a *horizontally integrated* business. It has been dominated by "integrated device manufacturers" such as Intel, AMD, Motorola, NEC, Hitachi, and Texas Instruments. These companies design, build, manufacture, and sell their own unique microprocessors. This vertical organization is evolving toward a horizontal one that supports just-in-time chip making. The new layering separates developers of intellectual property "cores," system-on-chip (SoC) designers, and "foundries." Profits will migrate to the foundries—outsourced semiconductor manufacturing—and to the developers and distributors of intellectual property cores (modularized circuits). Companies in the new layering include ARM (ARMHY) and Tensilica, which offer microprocessor cores, and TSMC (TSM) and UMC (UMC), which are foundries.

build the chips, and organizations for sales and distribution. They do the whole thing, from idea to silicon.

There's pressure on these manufacturers to add new instructions and new features because the general-purpose microprocessor isn't a perfect match for a particular customer. Large customers twist the manufacturer's arm to get instructions and features that benefit that customer's application. Manufacturers assimilate customer requests to plan the next-generation microprocessor. Adding upward-compatible instructions and features (instructions and features that grow from a compatible base) requires extending the development tools. New instructions have value only if the software development tools—compilers, assemblers, and debuggers—know about them.

The interval between the customers' requests and fielding a new microprocessor (with software tools) can be *years*. Further, since the manufacturer makes design tradeoffs across its installed base, it is unlikely that the customer will get all the desired features. Plus, as an off-the-shelf part, new features are available to the customer's competitors.

Designing products using off-the-shelf microprocessors makes it difficult to gain competitive advantage in the product itself. No off-the-shelf microprocessor fits a customer's application perfectly. A company might, therefore, be tempted to design a custom microprocessor for its application. The microprocessor's features wouldn't be available to competitors, giving the developer an advantage. But this means a delay of years to develop and to field the microprocessor. The company will have to bear the cost of the microprocessor's design and the cost of custom compilers, assemblers, and debuggers. Custom microprocessor development may cost several hundred million dollars. In-house, custom design offers competitive advantages and it offers excellent performance, but it takes too long and it costs too much.

Embedded x86. Think x86 and you think "PC," but, pretty soon, the x86 will invade the embedded market. The x86-based PC created the programs, data, and file formats that constitute the World Wide Web. As devices connect to the Internet, they will be drawn to x86 for compatibility. The few x86-based microcontrollers will rise in popularity. In September 2000, National Semiconductor (www.national.com) announced three x86-based microcontrollers designed specifically for Internet-connected applications. The Geode SC1200, SC2200, and SC3200 are optimized for set-top boxes, for thin-client desktop computing, and for portable devices, respectively. At the Microprocessor Forum last October, National announced the Geode GX2 (www.national.com/gx2) improved x86 microcontroller for shipment in the first half of this year. The GX2 crams the PC's CPU and most of its subsystems and I/O into a chip that dissipates less than one watt. STMicroelectronics offers a family of x86-based microcontrollers (www.st.com/stpc) and Transmeta has announced an x86-based microcontroller, the TM6000, for shipment in the second half of this year.

The next generation of mobile Internet devices won't immediately convert from ARM to x86. Take cell phones, there's a huge base of software already written for the ARM and not for the x86. The likely transition will be to have the legacy software running alongside an x86 core. The x86 core enhances compatibility with connected applications. Over time, the x86 will stay and the applications that once ran on the original host will migrate to the x86. If you develop software for cell phones you have to port it to all the processors in use with each update. If you know there will be an x86 in every cell phone, you will save yourself some effort by porting the software to the x86 and ignoring the other dozen processors. Three processors (legacy, x86, DSP) in an untethered device is a luxury that cannot be sustained for long, so there will be pressure to lose one or two. There aren't enough skilled engineers. Anything that makes work easier or minimizes engineering hours will be successful.

Two developments are changing the embedded-microprocessor market: the increasing popularity of microprocessor cores and the increasing sophistication of software tools. More on this later.

Microcontrollers

The microcontroller is a conundrum. One of the microprocessor's principal advantages is that it is a general-

DynamicSilicon

Editors	**Nick Tredennick**
	Brion Shimamoto
Publisher	**Lauryn Franzoni**
Designer	**Julie Ward**
Subscription Services	**Melissa McNally**
Editorial Director	**Richard Vigilante**
President	**Mark T. Ziebarth**
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.

purpose device that can be applied across a range of applications. The range of applications increases unit volumes, which lowers cost. The microcontroller combines the general-purpose microprocessor with peripheral functions specific to a narrow range of applications. Adding peripheral functions to the chip negates the microprocessor's production-volume advantage to gain efficiency in specific applications. Yet that's where today's microcontroller market is. Integrated device manufacturers, such as Intel, Motorola, Texas Instruments, Fujitsu, Hitachi, and Matsushita design, manufacture, and sell microcontrollers. There are thousands of microcontrollers, each aimed at a particular market segment. Unique microcontrollers fragment production, which increases chip cost. Because there are so many unique microcontrollers, it's hard for embedded system designers to find the right one for their application. It's likely that no microcontroller will have exactly the right set of features to fit the application.

Most microcontrollers aren't built in expensive, leading-edge processes. They don't have to be. Microcontrollers are designed for low cost and for adequate performance—they only have to be good enough to do the job. The microcontroller doesn't need 200 million transistors and a 2-GHz clock to make decisions in a blender or a hair dryer. Most microcontrollers have fewer than a million transistors and run at less than 30 MHz. Once all the transistors for the job fit on the smallest practical chip, making the transistors smaller with a more expensive process only increases cost.

Triscend, Hitachi, and Cypress Microsystems (a subsidiary of Cypress Semiconductor (CY)) take advantage of programmable logic to reverse the proliferation of microcontrollers. Instead of continuing to fragment the microcontroller market with more custom chips to meet diverse requirements, these companies build chips that adjust to diverse requirements.

Triscend chips, for example, offer a microprocessor core and a standard set of peripherals plus a block of programmable logic. Using Triscend's development software, the engineer chooses from a library the set of peripherals that fits the application. The engineer gets a custom microcontroller suited to the application, but Triscend's foundry manufactures a common chip that is "personalized" after it leaves the foundry. The common chip consolidates production across applications and the application gets a better-tailored microcontroller.

CPU microprocessors

Since the introduction of the IBM PC, CPU microprocessors have been on a quest for performance. For PCs, performance *and* cost are important. For a time, microprocessor-based workstations, from companies such as DEC, Sun, MIPS, SGI, Intergraph, and HP, contended for a share of the PC market.

PC makers built for *volume*; Moore's law supplied performance. Workstation makers built for performance; nobody supplied volume. In the competition between CPU microprocessors and workstation microprocessors, the volume strategy beat the performance strategy. It is ironic that Intel, the winner of the CPU microprocessor *volume* strategy, is repeating the mistake of the workstation microprocessors with its Itanium microprocessor. Itanium is built to enter the high end of the microprocessor market in competition with high-end x86-based microprocessors. Apparently, Intel's Itanium strategy is to enter the high end of the microprocessor market and then to move down. Meanwhile, its own x86 microprocessors, entering the market at lower performance but at much higher volumes, will evolve faster and will overtake Itanium implementations in their own market segments.

While Intel and AMD own the CPU microprocessor market, they are trapped by that market. Intel and AMD run a treadmill of ever-increasing performance to support dwindling demand for greater PC performance. They do this because that's where the margins are today. Intel and AMD design for the high-performance market segment. They believe they can serve other market segments with *derivative* products (made by modifying older generations of performance-oriented microprocessors). This is a strategy that assumes performance still drives the market. Modified performance-oriented microprocessors aren't optimized for energy-efficient applications, such as laptops and notebooks. The Formula 1 race car cannot become a Honda Insight.

Transmeta and VIA see the market shift from performance to energy efficiency and to price-performance. Transmeta and VIA design x86 microprocessors that conserve power and chip area (cost). Transmeta had trouble delivering chips late last year as it changed foundries from IBM to TSMC. That difficulty is behind it and Transmeta is ramping production in TSMC's 0.13-micron process. Transmeta has design wins in blade servers, in laptops, and in small PCs. VIA's products, because they are designed specifically for this segment, are a better match for Asia's emerging value-PC market than the derivatives of performance microprocessors from Intel and AMD.

As microprocessor performance rises, more users' expectations are satisfied by less-than-leading-edge microprocessors. In the early days of the PC market, the user base eagerly awaited each higher-performance generation.

Today, "value PCs," which use microprocessors a generation or two behind the leading edge, are probably the fastest growing market segment.

The PC's emphasis is changing from performance to energy efficiency.

Energy-efficient servers. Rising Internet use increases demand for files and web pages. Files and web pages come from server farms, which are rooms full of rack-mounted "server" PCs. Increasing demand means more PCs. More PCs means more floor space and more power. Floor space is expensive; power is worse—power could be $1000 per month per server rack. For each watt of PC power, add six-tenths of a watt in cooling. But the real problem is that the server farm cannot grow beyond the power limits of the building's wiring or of the power company's local distribution grid.

The market for small, energy-efficient servers will grow rapidly over the next few years. Seven times as many of these new servers fit in the same floor space. And they use one-fifth the power of today's servers. Gartner Dataquest expects the "blade server" market to grow from fewer than one hundred thousand this year, to more than a million units per year by 2005. An International Data Corp. estimate puts the number at two million and estimates revenue at $2.9 billion by 2005.

Startups, such as RLX Technologies (*Dynamic Silicon*, Vol. I, No. 11), Egenera, FiberCycle, and Rebel.com, recognized the opportunity and either built or announced energy-efficient, high-density servers. RLX Technologies, for example, builds "blade" servers, which squeeze 324 servers into a rack that usually houses 42. In strategic moves to reserve the market, HP, IBM, Dell, and Compaq announced blade-server products. Announcements by large companies killed FiberCycle and Rebel.com last fall. IBM partnered with RLX Technologies and will sell its products.

RLX Technologies builds blade servers around TM5600 microprocessors from Transmeta. Transmeta's microprocessors are optimized for low-power applications. They are particularly well-suited for blade-server applications because they run cool enough to be placed on the board without a finned heat sink or fan, which means smaller servers and therefore more servers per square foot.

Digital signal processors

Digital signal processors (DSPs) are great for data- and computing-intensive calculations, but they still use a lot of power. Using more power for ease of design or for lower cost is OK, as long as the system plugs into a wall socket. But the world is splitting into tethered and untethered devices. Tethered devices, such as DVD players, televisions, and VCRs plug into wall sockets for power and will continue to employ digital signal processors. Untethered devices, such as cell phones, MP3 players, and GPS receivers, use digital signal processors today. To conserve power, digital signal processors will evolve from being standalone chips to being soft cores.

DSP unit shipments are growing rapidly because DSPs do what they do—data-intensive computations—better than general-purpose microprocessors. Integrated device manufacturers such as Texas Instruments, Agere, Motorola, and Analog Devices (ADI), make DSPs. Integrated device manufacturers sell chips; but in the future, their customers will want cores. There's pressure to reduce the number of components in the cell phone to make the battery last longer and to make the cell phone cheaper (fewer chips) and more reliable (fewer chip-to-chip connections). Integrating the DSP with the microprocessor and with other functions as a system-on-chip reduces the number of components.

Today's situation makes the transition from standalone chips to cores and to system-on-chip design difficult. Existing DSP software was developed for standalone chips from the integrated device manufacturers, so the instruction sets for the standalone chips are entrenched. DSP cores are available from companies such as 3DSP, Lexra, Bops, DSP Group, and Infineon. These DSP cores, however, are not compatible with the legacy software, so it will be difficult for cores to get design wins in the highest-volume, program-intensive applications (such as cell phones, MP3 players, and GPS receivers). Texas Instruments, the market leader, offers cores to customers of its system-on-chip designs, but these cores are not portable to foundries.

Vendors of soft embedded microprocessor cores, such as ARC, ARM, and Tensilica are adding digital signal processing extensions to their cores. These extensions to entrenched microprocessors provide the bridge for legacy DSP software to move from the standalone DSP to the microprocessor core.

Cores

As more transistors fit on a single chip, standalone microprocessors become less valuable (seen in fig. 1 as declining share for embedded microprocessors against microcontrollers, which are more integrated). Chip-to-chip communication wastes board space, power, and time. Better to have the processor on the chip with the rest of the functions. Even a sophisticated 32-bit processor occupies only a fraction of the chip. More than 600 copies of Tensilica's base soft-core microprocessor would fit on a chip the size of Intel's Pentium 4. Designers collect logic blocks and intellectual-property cores and assemble them on a single custom chip called a system-on-chip. This has given rise to system-on-chip suppliers such as Fujitsu, IBM, and LSI Logic and to suppliers of microprocessor cores such as ARM, MIPS, Bops, Nazomi Communications, TriMedia, and Lexra.

The rise of microprocessor cores changes the market. Integrated device manufacturers want to supply chips while more of their customers want cores. Since the customer assembles a collection of cores (a microprocessor core and other types of cores), it is unlikely that all cores will come from a single source, so the customer wants "soft" cores. Soft cores are modularized circuit descriptions not tied to one semiconductor process. Soft cores from diverse sources can form a system-on-chip design that can go to any foundry. An integrated device manufacturer will be reluctant to supply soft cores to customers because it wants customers that use its cores to use its production plant (fabrication plants are expensive, so it's good to keep them full). A "hard" core is a modularized *physical* circuit; it assumes a *particular* semiconductor process. Hard cores can be smaller and more efficient than equivalent-function soft cores, but they are not portable among foundries.

ARC and Tensilica. ARC Cores (ARK) and Tensilica offer flexible microprocessor cores along with flexible software development tools. The general-purpose core is complete enough to implement ordinary control functions. Developers add new instructions to tailor the microprocessor to a particular application. Sophisticated software from ARC and from Tensilica grasps the new instructions and spits out support in the form of compilers, assemblers, debuggers, and operating-system extensions.

Custom instructions can improve application performance by factors of 10 to 100 over the application's performance on a general-purpose microprocessor. Applications include cell phones, portable digital assistants, routers, telecommunications equipment, digital cameras, and digital video recorders. The increase in performance comes with savings in power (fewer, more efficient instructions), with lower cost (one component may replace several), and with new features (unique instructions). The developer gains the advantages of custom instructions without the cost or delay of developing a custom microprocessor and its software tools. Great tools make it possible to develop a custom microprocessor in hours.

Designing a custom version of an ARC or of a Tensilica microprocessor doesn't require the expertise of a microprocessor designer. Knowing *what* features or instructions to add is more important than knowing *how* to add them. ARC and Tensilica automate the adding of features and instructions. This increases the pool of designers capable of developing a custom microprocessor from a few thousand to tens of thousands. Enabling more designers means a larger market opportunity for ARC and for Tensilica.

The next generation of software from Tensilica makes the process even simpler. Instead of specifying custom instructions, the engineer specifies the total behavior of the microprocessor and of the program running on it. Tensilica's software searches alternatives and builds custom instructions *and the application code* for the custom microprocessor. With this generation of custom microprocessor development software, hundreds of thousands of embedded-systems designers will be able to design custom microprocessor cores for specific applications.

ARC and Tensilica make designing custom microprocessors for embedded applications simpler. By making custom microprocessor design simpler, they have increased the audience of potential customers from thousands to tens of thousands and soon their software will enable hundreds of thousands of engineers. As ARC and Tensilica gain market share, the introduction of new electronic products will accelerate as the limited resource of development engineers becomes more productive.

ARM Holdings, Ltd. ARM is the 800-pound gorilla of hard and soft embedded microprocessor cores, thanks to its adoption in cell phones. ARM's dominance in cell phones becomes entrenched as more application software is developed. ARM is entrenched in cell phones in the same way that Intel is entrenched in PCs. New entrants face increasing difficulty in persuading handset manufacturers to port applications to non-ARM microprocessors. Despite the downturn in electronics, ARM, which gets 50% of its royalties from cell phone makers that use ARM microprocessor cores, posted impressive gains. ARM's 2001 pretax profits rose 42% to $71 million. Sales for 2001 were up 9% from 2000 to $206.3 million (the rest of the industry *shrunk* by 32% in 2001).

Because of its early start, ARM is defining the intellectual property core business. It has moved from hard to soft embedded microprocessor cores and it is now moving to soft microcontroller cores (which include peripherals and digital signal processing extensions).

An x86 core. The x86-based microcontroller offerings from National, STMicroelectronics, and Transmeta are steps in the right direction, but they aren't enough. As the world splits into tethered and untethered devices, there is pressure to move from physical chips to cores for system-on-chip implementations for the power-sensitive untethered applications. ARC, ARM, Tensilica, and a host of other companies offer microprocessor cores for system-on-chip applications, but there's no x86 core. Transmeta and VIA Technologies (www.via.com.tw) are in the best position to offer x86-based microprocessor cores since both are fabless and both have Pentium-class designs. I'd like to see a soft x86 microprocessor core on a roadmap from one of these companies.

Foundries (TSMC, UMC, Chartered)

Semiconductor fabrication plants cost billions of dollars. Buying the equipment isn't the only huge cost. Developing the process is enormously expensive. (The distinction between equipment and process is the distinction between the appliances in a kitchen and the recipes and procedures used in cooking.) It was once an advantage for integrated device manufacturers to own plants and processes. In the early days of foundries, integrated device manufacturers had better processes than the foundries, so the integrated device manufacturers had a competitive advantage.

Since they didn't have leading-edge processes, early foundries built chips for companies developing new markets. In an emerging market, a less-than-leading-edge chip is better than no chip at all. Early foundries provided chips to Altera and to Xilinx as the programmable logic market grew. The foundries improved their semiconductor processes to maintain a competitive edge for their fastest-growing customers in each new market. Programmable logic devices for Altera and for Xilinx and graphics chips for Nvidia and for ATi Technologies drive process improvements in today's foundries. The foundries have caught up and the tables have turned.

Small design teams have access to the leading-edge semiconductor processes and to the intellectual property libraries of the foundries. The semiconductor processes are as good and the intellectual property libraries are quickly becoming better. These changes erode the integrated device manufacturer's competitive advantage. Further, if an integrated device manufacturer builds production capacity to meet peak demand then expensive excess capacity will have to be sustained through non-peak production. Integrated device manufacturers have begun to partner with foundries both for semiconductor process development and for production.

Today's leading-edge process is 130 nanometers (0.13 microns); the next-generation process is 100 nanometers (0.10 microns). How can process development be so expensive? It doesn't sound too difficult; just shrink the 130-nanometer features to 100 nanometers. Process development isn't that simple. The design rules for a particular process may run hundreds of pages. *Think hundred-page recipes*. Every detail of line width, corners, spacing between lines, thickness, aspect ratio, proximity, and chemical composition is important. And they are interdependent. Change an innocuous feature in one rule and problems appear in another area. Getting each generation of design rules right is an expensive effort.

NEC, LSI Logic, and Conexant have partnered with TSMC. AMD, Hitachi, IBM, and Infineon have partnered with UMC. Motorola and Lucent have partnered with Chartered (CHRT). Even Intel and Texas Instruments have partnered with foundries. Texas Instruments already outsources 5% of its production to foundries and may increase outsourcing to as much as 20%. It makes sense. Contracting with a foundry to handle excess production leads the integrated device manufacturer to work with the foundry on process development (transferring production is simpler if the processes are compatible). Cooperation on process development shares cost among participants, which makes the entire industry more efficient, which spurs growth.

Building chip production capacity is expensive and it takes a couple of years. A company may have to build capacity in a downturn in order to have it online for the next growth cycle. Building capacity in a downturn based on two-year market forecasts is risky. It's smarter for integrated device manufacturers to build more conservatively and to contract with foundries for excess production. This shifts production-capacity risk to the foundries. Because they have a varied customer base, the foundries can manage the demand for capacity better. As a consequence of sharing both process development and production, semiconductor processes become more standardized across the industry, making designs more portable.

Lessons

Embedded microprocessors, microcontrollers, CPU microprocessors, and digital signal processors developed in an environment of vertically integrated device manufacturers. New opportunities open as the industry shifts from the vertically integrated companies with one-stop shopping for design, manufacture, and sales to horizontally integrated companies of intellectual property cores, system-on-chip designers, and semiconductor foundries.

Today's embedded microprocessor market, characterized by proliferation of instruction sets and by long lead times for new chips, is changing as users demand microprocessor cores and rapid customization. The emerging leader in microprocessor cores is ARM; the emerging leaders in rapid customization are ARC and Tensilica. The intellectual property business is more difficult than it seems. Income is typically based on a fixed fee plus royalties. One-and-a-half to two years may elapse between licensing and royalty income—and that's only if the project isn't canceled or derailed. Even if the project is completed, royalties won't be significant unless the product sells well. Large companies are reluctant to license from startups and small companies can't afford hefty fixed fees.

Component proliferation characterizes today's microcontroller market. Triscend, Cypress Microsystems, and Hitachi will consolidate the microcontroller market with microcontrollers that offer software-selected peripheral functions that designers configure in the field.

Intel dominates today's CPU microprocessor market. The CPU microprocessor's performance orientation is changing to power conservation. The x86 microprocessor market's orientation will change from performance-based desktop PCs to low-power small PCs and to x86-based microcontrollers. The transition to small PCs and to x86 microcontrollers favors Transmeta and VIA Technologies. The next transition, to x86 microprocessor cores, also favors Transmeta and VIA Technologies.

Digital signal processors have grown fastest among microprocessor segments. The transition will be slow, but chips will give way to digital signal processing extensions to microprocessor cores from companies such as ARM, ARC, MIPS, and Tensilica.

Integrated device manufacturers, such as NEC, Hitachi, Motorola, Intel, and Texas Instruments, built huge businesses in embedded microprocessors, microcontrollers, CPU microprocessors, and digital signal processors. These companies are vertically integrated; they design, build, and sell chips. The world is changing to favor a horizontal integration consisting of intellectual-property cores, system-on-chip designers, and foundries. Foundries, such as TSMC, UMC, and Chartered will grow rapidly in this new environment.

Integrated device manufacturers will struggle with the transition. For example, if I ran Intel, the largest integrated device manufacturer, here's what I would do. Dump non-x86 microprocessors (Itanium). Continue to build high-end x86 microprocessors for PCs. Build x86 microcontrollers and cores, concentrating on low-power implementations. License *all* products as soft cores. Align with at least one foundry on semiconductor process and sell wafer starts on the open market. Make Intel's internal organization parallel the external organization of intellectual property cores, system-on-chip designers, and foundries.

Nick Tredennick and Brion Shimamoto
February 14, 2002

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Altera, Xilinx	Fabless	Excellent	Fabless programmable logic leaders in a growing market. Already positioned for the future.
ARM	Fabless	Excellent	Early to market with microprocessor core licensing. Already moving to soft microprocessor and microcontroller cores.
Tensilica	Fabless	Excellent	Leader in configurable soft microprocessor cores. Excellent strategic vision.
TSMC, UMC	Foundry	Excellent	Top two fabless semiconductor foundries. Cooperative agreements with many integrated device manufacturers.
ARC	Fabless	Good	Leader in configurable soft microprocessor cores. Uncertain strategic leadership.
Chartered	Foundry	Good	Number three fabless semiconductor foundry. Position slightly weaker than top two foundries.
Cypress Microsystems, Triscend	Fabless	Good	Microcontrollers with drag-and-drop peripherals in programmable logic. Could consolidate microcontroller market.
Transmeta, VIA Technologies	Fabless	Good	x86 microprocessors for energy-efficient and for portable applications. Positioned to offer x86 cores.
Analog Devices	Integrated	OK	Broadly based in analog, ASSPs, and MEMS. Business tolerates slow transition.
IBM	Integrated	OK	Broadly based in systems and components. Experience with vertical to horizontal transition in mainframe systems.
National	Integrated	OK	Broadly based in consumer electronics. Good position with x86 microcontroller. Business tolerates slow transition.
RLX Technologies, Saint Song Corp.	Systems	OK	System houses not strongly affected by vertical to horizontal transition in semiconductor manufacturing. x86-based systems.
Texas Instruments	Integrated	OK	Broadly based in analog, ASSPs, and MEMS. Business tolerates slow transition.
AMD, Intel	Integrated	Struggle	Entrenched chip culture and performance-oriented design. Difficult transition for prototypical integrated device manufacturers.
Fujitsu, Hitachi, Matsushita, Motorola, NEC	Integrated	Struggle	Microcontrollers. Difficult transition from chips to cores and from in-house manufacturing to selling wafer starts on the open market.
Lexra, MIPS	Fabless	Struggle	Microprocessor cores, but not entrenched in high-volume, long-term systems.
LSI Logic	Integrated	Struggle	Difficult transition from hard cores for ASICs to licensing soft cores and from in-house manufacturing to selling wafer starts on the open market.
STMicroelectronics	Integrated	Struggle	Broadly based in microprocessors and components, including x86, but unclear strategy.
3DSP, Bops, DSP Group	Fabless	Fail	DSP functions will migrate from standalone chips to embedded microprocessor core extensions, bypassing DSP cores.
Agere	Integrated	Fail	Culture of regulated business doesn't translate well to competitive environment.

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	1/31/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	25.12	14.66 - 33.60	9.68B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	43.80	29.00 - 64.00	16B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.59	£0.25 - 3.17	£109M
ARM Limited (ARMHY***)	Microprocessor and System-On-A-Chip Cores	11/26/01	16.59	14.35	8.39 - 23.63	4.92B
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	3.16	1.48 - 11.48	84M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	26.44	16.06 - 37.13	3.6B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	21.76	13.72 - 29.25	2.64B
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††	16.97	8.39 - 20.14	57.13B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	3.14	1.17 - 37.25	422.8M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	8.82	4.25 - 10.45	23.39B
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	18.02	9.71 - 38.5	1.4B
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	43.35	19.52 - 59.25	14.48B

† 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.

Company Update

"While Altera and Xilinx diligently position themselves for prototyping and for displacing ASICs, they have been overlooking a huge untapped market: positioning themselves against DSPs and microprocessors." *Dynamic Silicon*, March, 2001.

On February 11, 2002, Altera announced its Stratix family of programmable logic devices (PLDs). These devices, the industry's largest and fastest PLDs, are optimized for DSP applications, for computing-intensive applications, and for system-on-chip designs. In addition to the usual programmable logic, the new chips offer dedicated DSP blocks, millions of memory bits, and high-performance interface logic. The company will distribute sample parts in the second quarter.