Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 7 | The Investor's Guide to Breakthrough Micro Devices | July 2002

Systems-on-Chip to the Fore

Many *Dynamic Silicon* companies sell microprocessors. My list includes Analog Devices, ARC, ARM, Cypress, National Semiconductor, Tensilica, Transmeta, and Triscend. You might even count Altera and Xilinx, since some of their chips contain microprocessors. These companies have different approaches to microprocessors. The diversity of suppliers is the artifact of an industry in transition. Microprocessors, along with all other types of chips, will be sucked onto super chips called "systems-on-chip" (abbreviated "SoC"). Chips will not be designed as complete physical entities, but as collections of smaller units of intellectual property called *soft cores*. The new units of design—soft cores—are logical descriptions that are independent of the chip manufacturing process. Soft cores designed by one company can be put on the same chip with soft cores designed by others. This allows the semiconductor industry to stratify horizontally into companies that design soft cores, into companies that pull together—aggregate—cores, and into companies that manufacture super chips. This new industry layering of component designs (soft cores), SoC designs, and SoC manufacturing is under way. Today, for example, ARC and ARM design soft cores and TSMC and UMC manufacture SoCs. Many companies aggregate soft cores.

Microprocessors do two things. The first is technical: microprocessors are "state sequencers"—they decide what to do next (e.g., three buttons got pushed, which one has priority?). The second is economic: microprocessors allow engineers to implement product functions in software, because performing functions in software is cheaper than building the functions in special hardware, and the microprocessors are fast enough. The criterion has always been cost-performance but this metric is about to be displaced by a new one, power-efficiency. Fewer watts mean longer-lasting batteries. Increasingly, our electronics will not be plugged into the wall. Microprocessors providing product functions through software are not power efficient. Systems-on-Chip will displace microprocessors as physical building blocks.

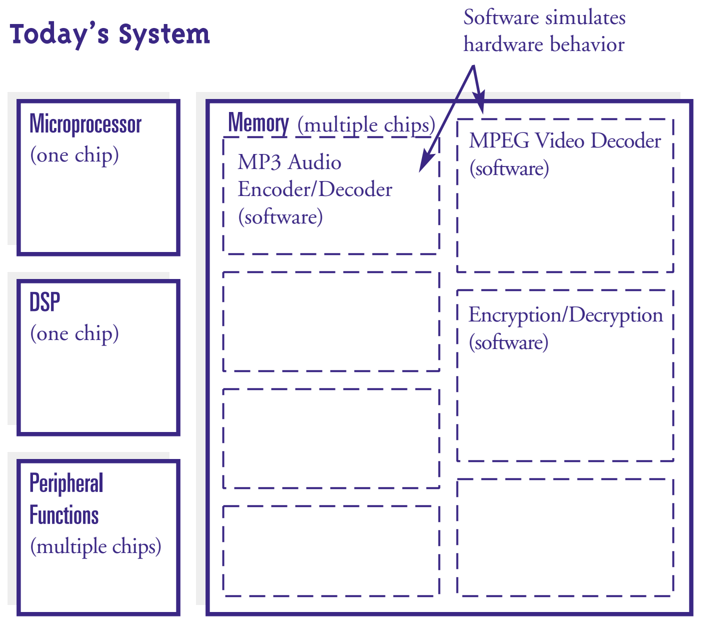

Fig. 1. Software running on the microprocessor and DSP simulates functions that would otherwise cost more in hardware.

In This Issue:

From a manufacturing view, chips can be divided into two types: generic and specialized.These are relative terms, but generic chips include memory and programmable logic devices (PLDs). Compared to these, other chips: microprocessors, DSPs, microcontrollers, ASICs, and ASSPs, are specialized. Generic chips have a large number of uses so the manufacturing cost is spread across a big number. Idea: why not express the design of specialized chips in a way that it can be put on a PLD? Personalize the generic chip, after manufacturing, with the circuits of the specialized chip? Bingo, "system-on-chip." Microprocessors recede, systems-on-chip begin their move to the forefront.

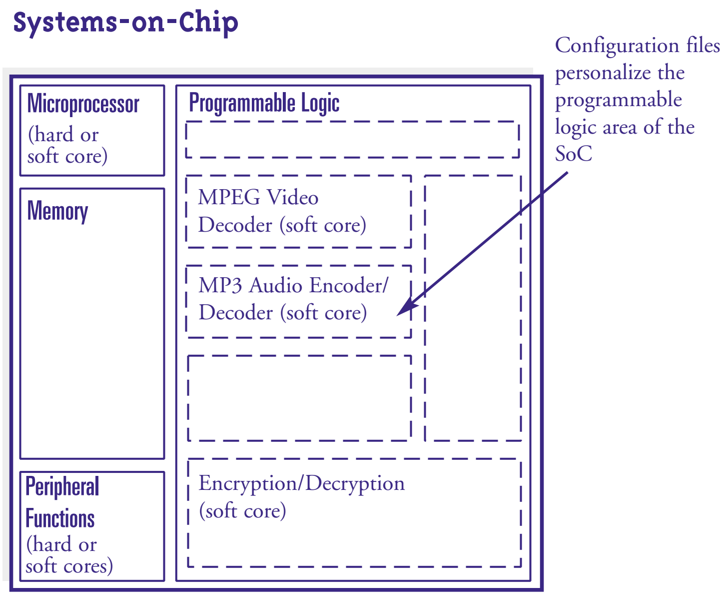

Fig. 2. As power efficiency becomes more important, functions implemented in software will become *soft cores* implemented in programmable logic.

Systems-on-chip satisfy two needs: 1) tailored hardware for power efficiency, and 2) generic hardware for high-volume manufacturing. Programmable logic allows systems-on-chip to become generic hardware. Programmable logic is conceptually two levels. One level is a sea of logic elements and wires. The second level is a "personalization memory." Loading the personalization memory with a string of ones and zeroes causes *physical connections* to be made between selected logic elements and wires. Physical circuits from ones and zeroes!

DynamicSilicon

Editors	**Nick Tredennick**
	Brion Shimamoto
Publisher	**Lauryn Franzoni**
Designer	**Julie Ward**
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.

Implications of the change to SoCs:

■ Designs move from being physical chips to being soft cores that can be placed on any chip.

■ DSPs go away as such. Their function is done in programmable logic.

■ Microprocessors remain in the form of soft cores for their what-to-do-next decision-making ability.

■ ASICs, ASSPs, and microcontrollers are subsumed in SoC. (ASICs and ASSPs are complex chips designed for a few customers. Microcontrollers are in everything. Manufacturers ship billions of microcontrollers every year.)

Below, I trace microprocessor development to explain why there are so many microprocessor vendors among the Dynamic Silicon companies. Four characteristics differentiate microprocessor markets: microprocessor architecture, design objective, delivery form, and flexibility. A microprocessor's architecture includes its instruction set, which determines its compatibility with existing software. Its design objective orients a microprocessor toward performance, energy conservation, low cost, or some balance. A microprocessor can be delivered as a physical chip or as a design description. Finally, the microprocessor's flexibility to adapt to different uses varies widely. These characteristics differentiate the customers and the applications. For example, sales of power-efficient microprocessor cores do not compete with sales of physical chips for desktop PCs.

Microprocessor architecture

Architecture defines the programmer's model and it has been a hot topic since the computer was invented. The programmer's model defines the computer's instructions, its registers, its view of memory, and its operating modes—all the features of the computer's abstract problem-solving model. Should the computer use fewer long-running instructions or should it use many fast instructions? Should arithmetic instructions operate on data in memory or only in registers? Engineers and university professors devote careers pondering these issues.

An architecture is only as important as its installed base of software. Writing software introduces a long delay in time to market. It takes a long time to create the installed base, because writing software is expensive.

Unfortunately, microprocessor architecture is swept by fads in the way that the fashion and toy industries are. RISC (reduced—as in simple—instruction set computing), a fad that started in the early 1980s, promised twice the performance at half the cost, but didn't deliver. Twenty years of commercial, conference, and academic debate ensued. The industry has learned that it is the sophistication of the semiconductor manufacturing process, not the sophistication of the architecture, that matters for performance. In other words, selling enough volume to be able to afford the next improvement in chip *manufacturing* is what matters for performance.

Be wary of performance claims that are based on a new microprocessor architecture.

Today's instruction set design has two camps. One camp is x86-compatible microprocessors. The other camp is everything else. AMD, Intel, National Semiconductor (NSM), STMicroelectronics, Transmeta (TMTA), and Via Technologies are in the x86 camp. Everyone else is in the "other" camp. Sun's SPARC, MIPS' MIPS, and IBM's PowerPC began as RISC microprocessors. All have added complex instructions.

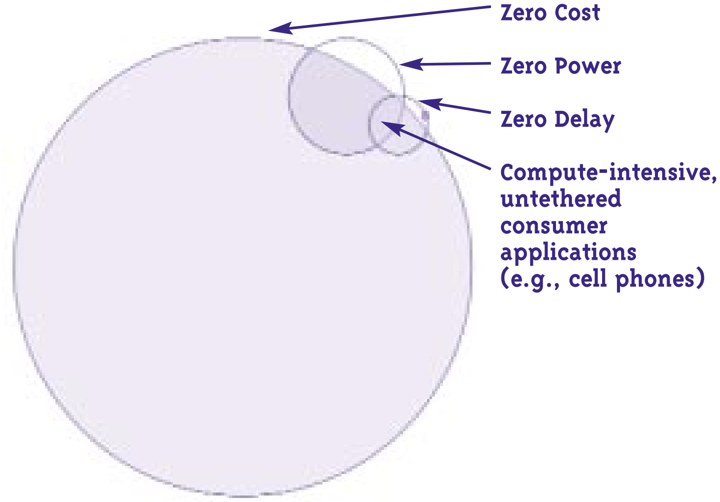

Fig. 3. The zeroes model for electronic systems market segments.

Design objective

I've talked before about semiconductor applications in terms of a "zeroes model" (*Dynamic Silicon*, Special Report, "MEMS and Dynamic Logic: Why Now?"). Fig. 3 illustrates the zeroes model. The overlapping circles represent microprocessor unit volumes. Zero cost represents the lowest-possible-cost consumer applications. Zero power represents the lowest-possible-power untethered applications. Zero delay represents performance-oriented applications. The overlap of zero delay, zero power, and zero cost is compute-intensive, untethered consumer applications. A similar model applies to microprocessor design objectives. Workstation, server, and PC microprocessors emphasize performance. Microprocessors for embedded consumer applications, such as blenders and electric razors, emphasize low cost. Microprocessors for compute-intensive, untethered consumer applications, such as cell phones, MP3 players, and GPS receivers, try to balance performance, power, and cost.

Delivery Form

Chips. Integrated device manufacturers (IDMs) delivered the first microprocessors as chips. These vertically-integrated companies, such as Intel, AMD, Motorola, and IBM, have their own architects, logic designers, circuit design and layout engineers, and manufacturing lines. More than thirty years after the microprocessor's introduction, IDMs selling physical chips still dominate microprocessor unit volumes.

The market is changing. Increasing transistor budgets bring the microprocessor's peripheral functions onto the chip, creating the microcontroller. Further integration creates application-specific integrated circuits (ASICs), application-specific standard products (ASSPs), and systems-on-chip (SoC) designs. An ASIC is a complex logic chip designed for a single customer. An ASSP is a complex logic chip designed to meet the needs of a range of customers.

Cores. Cores are modular designs meant to be placed onto a chip to complete a larger design.

Hard cores. Microprocessors designed for the first microcontrollers, ASSPs, and ASICs were delivered as "hard cores." A hard core is a circuit description that is closely tied to a particular semiconductor process. Changing the semiconductor process means redesigning the hard core. Moving to a new foundry means redesigning the hard core.

Soft cores. A "soft core" is a logical description that is *independent* of the semiconductor process. A hard core is more efficient and is likely to have better performance than a soft core, but the soft core is *portable* across semiconductor process generations and between foundries.

The IDMs will fragment horizontally because fabs—the physical manufacturing plants—are expensive and because developing the semiconductor process—materials and procedures—is expensive. The semiconductor industry is cyclical. If an IDM builds capacity to support good times, then expensive capacity goes idle in bad times. It's better to build less capacity and to contract with a foundry for extra production. Since the foundry supports a wider range of customers than the IDM supports, the foundry is better able to absorb production variations. Offloading work to a foundry implies cooperation (and beneficial cost-sharing) on process development. Designs become more portable; the IDM is more likely to buy cores it needs to build its ASSPs; it is more likely to license cores as well.

As the IDM model fragments horizontally, the dominance of chips will slowly give way to hard and soft microprocessor cores. During the transition, there are profitable niches in all three forms (chips, hard cores, soft cores).

Flexibility

The microprocessor's architecture, design objective, and delivery form help differentiate markets, but the biggest differentiator for Dynamic Silicon companies is flexibility. I divide microprocessor flexibility into three categories: *fixed*, *configurable*, and *reconfigurable*. Configurable microprocessors further subdivide according to whether the microprocessor's resources, its instructions, or its peripherals, are configurable.

Fixed. National Semiconductor and Transmeta sell x86 chips (rather than cores). For National and for Transmeta and for their customers, the instruction set is fixed. The x86 instruction set is set by Intel and by Microsoft. Intel and Microsoft can modify the instruction set, but modifications will remain "backward compatible." Backward compatible means that a new microprocessor will run old programs in their binary form. New x86 microprocessors must be able to run more than twenty year's worth of old programs.

It isn't just that the instruction set is fixed for x86 microprocessors. Compatibility with the finest detail of the instruction set is an absolute requirement. A slight deviation in one obscure case in a floating-point instruction renders the microprocessor useless. A fixed instruction set is the foundation of this microprocessor's market.

Configurable. A microprocessor is configurable if the customer specifies resources (e.g., caches, arithmetic units), instructions, or peripheral functions (e.g., serial ports, counters, timers).

Fixing the instruction set and letting the customer adjust the microprocessor's resources sets the performance and cost bar along a software-compatible range of microprocessors. ARM and MIPS offer cores with configurable resources. As MIPS says, "The highly configurable and synthesizable core allows designers to include only the features necessary to their application." MIPS even allows a user-defined coprocessor and user-defined instruction-set extensions. Configuring resources is one of the options offered by ARC International and by Tensilica, though these companies also let their customers mess with the instructions. Letting the customer tailor the microprocessor's resources to suit an application is more flexible than having the supplier guess the family of software-compatible microprocessors that will suit its customers.

Microprocessor and computer instruction sets evolve as the supplier learns more about customers' needs. The microprocessor's new bit-manipulation instructions may help dozens of customers with their applications, but aren't a perfect fit for any. The customer wants custom instructions that will run its encryption application ten times faster, but it doesn't want to give competitors access to these proprietary instructions. Custom instructions make the microprocessor more efficient.

Letting the customer mess with the instruction set seems the worst choice. Once the instruction set changes, compilers, development systems, test programs, operating systems, and application programs have to know about the instruction changes if the application is to take advantage of them. If the instruction set changes, who's to say whether the mess the customer creates will even be functional? That would be true if these companies offered their customers a clean slate. Companies offering configurable instruction sets don't give the customer free reign. The customer works from a functional base. This is the hard problem that ARC International and Tensilica have solved.

One advantage of letting the customer mess with the instruction set is that it prepares the supplier for the next step in development systems—accepting an executable specification from the customer. The supplier automates building custom instructions. The customer specifies the requirements; software from ARC or Tensilica searches the design space for a good solution. The customer specifies the algorithm and its desired performance. That's easier than specifying new instructions. ARC and Tensilica raise the level of abstraction enabling more designers and making them more productive.

Allowing the customer to specify peripheral functions means the supplier is offering a microcontroller or systems-on-chip and not a simple microprocessor. That's where the highest-volume markets will be. As more peripheral functions migrated onto the chip with the microprocessor, the variety of microcontrollers ballooned. Each high-volume application wanted its own microcontroller with just the right set of peripherals.

Semiconductor process progress has shrunk the microcontroller's peripheral functions and its microprocessor core to a small portion of the chip size necessary to accommodate the chip's pins. Once the chip reaches the size needed by its pins, the chip can't shrink. If the chip can't shrink, it doesn't get any cheaper as the circuit on the chip shrinks. The extra area is free. Instead of building many microcontrollers, each with its own set of peripheral functions, Cypress Microsystems and Triscend replace the custom collection of peripherals with programmable logic. Accompanying software enables the customer to “drag and drop” peripheral functions to build a microcontroller that best suits the application. Cypress Microsystems and Triscend build a few general-purpose microcontrollers that are personalized by the customer to suit *thousands* of applications.

Reconfigurable. Configurable microprocessors allow the user to customize the hardware, the instruction set, or the peripherals before building the chip. Reconfigurable microprocessors move this flexibility from design to the application.

Instead of fixed decoders and execution resources, imagine the microprocessor as a sea of programmable logic. Want to run x86 programs? Load the x86 configuration and run the programs. Need a special instruction that uses 57 multipliers for three cycles? Configure the logic and the decoders and run the instruction. Configurations can be pipelined just like instructions (it’s all just logic), so there’s no sacrifice of latency or throughput. Ascenium, Proceler, and other startups are working on ideas such as these that offer *run-time* flexibility. This flexibility increases performance and it improves efficiency. The cost is the overhead for configuration transistors, but transistors are always getting cheaper. The largest problem for these startups is educating engineers and investors.

Untethered systems

The first integrated circuits were small logic functions (e.g., a few NAND gates or inverters). As the semiconductor process improved, the complexity and variety of these IC macro functions grew. Engineers built the state sequencer, computing resources such as registers and arithmetic units, and input and output logic, using chips from a catalog of IC macro functions. The state sequencer is the engine that controls a digital system. It initializes the system when the power comes on and it makes decisions about what to do next based on its history and on its present inputs. IC macro complexity grew until a single chip contained a state sequencer and some computing resources. It was the first microprocessor and it forever relieved engineers of designing state sequencers.

The microprocessor brought the computer’s programming model to embedded systems. The microprocessor provided not only the state sequencer, but it could also simulate functions that would otherwise be expensive hardware. Simulating hardware functions with a microprocessor compromised efficiency and performance, but it made the system cheaper by reducing the number of chips. Performance-critical functions stayed in hardware and the rest of the functions moved to the microprocessor. The microprocessor has had thirty years of simulating hardware functions in embedded systems. As the microprocessor’s performance improves, fewer performance-critical functions remain in hardware.

State and federal emission regulations changed the rules in automotive engine design. Engine efficiency became the most important design constraint. Similarly, new emphasis on untethered systems changes the rules in microprocessor design. Power efficiency becomes as important as performance and cost.

For systems that plugged into a wall socket, the level of performance needed determined whether a function was implemented in hardware or was simulated by the microprocessor. For systems with access to unlimited power, efficiency wasn’t important; for untethered systems, efficiency is important. As efficiency becomes more important, functions migrate from being simulated to hardware. In order to migrate to hardware without giving up the microprocessor’s flexibility, however, these functions will migrate to hardware using programmable logic devices. The microprocessor won’t disappear; its function as the state sequencer for the system remains.

Lessons

There are two types of microprocessor applications: those that have a graphical user interface and those that don’t. If the application has a graphical user interface, it will eventually be an x86. If the application doesn’t have a graphical user interface, it can use any instruction set. Abundant legacy software supports camera, cell phone, and set-top box applications, however, so the transition from today’s instruction sets to x86 will take years.

There are three market segments for stand-alone chips and two segments for microprocessor cores.

There’s a market for x86 chips for PCs and for other applications. The market for PCs will grow, but PCs have reached diminishing returns on performance. PCs have enough performance for most users. Early adopters have their PCs, so growth in the market will be late adopters and it will be in growing markets such as Asia. These markets favor “value” PCs based on mid-range and low-end microprocessors. Based on its opportunity in these emerging markets for x86 microprocessors, VIA Technologies (www.via.com.tw) is a Dynamic Silicon

Companies	Microprocessor Market Segments					Microprocessor Characteristics				
	x86 Microprocessors	x86 Microcontrollers	Non-x86 Microcontrollers	x86 Cores (no one!)	Non-x86 Cores	Configurable Resources	Configurable Instructions	Configurable Peripherals	Reconfigurable	In-house Fabrication
Altera*			■		■	■	■	■		
ARC International*			■		■	■	■	■		
ARM Holdings, Inc.*			■		■	■		■		
Cypress Microsystems*			■		■			■		■
National Semiconductor*		■	■							■
Tensilica*			■		■	■	■	■		
Transmeta*	■									
Triscend*			■		■			■		
VIA Technologies*	■									
Xilinx*			■		■	■	■	■		
AMD	■	■								■
IBM			■		■					■
Intel	■	■	■							■
MIPS Technologies			■		■	■				
Motorola			■							■
STMicroelectronics		■	■							■
Texas Instruments			■							■

* Dynamic Silicon Company

Fig. 4. Market segments, microprocessor characteristics, companies, and offerings.

company. VIA is a full-range supplier of chips sets and system boards to the value-PC market. Their products include x86 microprocessors and chip sets forgraphics, Ethernet, USB, Firewire, and CD-ROM/DVD. Whereas Intel's approach with OEMs might be characterized as "Here's the way you're going to do it."

VIA's approach is lighter, being "Here's what's possible--and it's cheap."VIA, in Taiwan, is well positioned geographically, technically and culturally to serve the Asian market. AMD and Intel will struggle as the market shifts toward cheaper chips.

There's an emerging market for x86 microcontrollers in appliances with graphical user interfaces and web connections, such as set-top boxes, personal digital assistants, and game platforms. This emerging market favors National Semiconductor and STMicroelectronics. Transmeta has also announced an x86-based microcontroller, so it too could benefit.

There's a market for non-x86 microcontrollers as stand-alone chips. These microcontrollers go into a wide variety of applications where there is need for a state sequencer and for peripheral functions, but there is no need for a graphical user interface. Microcontrollers are ripe for consolidation. There are so many instruction sets and peripheral combinations that it's difficult for an engineer to find the microcontroller that's perfect for a particular application. Cypress Microsystems, Hitachi, and Triscend have an opportunity to consolidate the microcontroller market by offering chips with a microprocessor core, some common peripherals, and a block of programmable logic. The microcontroller's corresponding development system lets the engineer "drag and drop" peripherals to build the perfect microcontroller for a particular application.

There's a nascent market for x86 microprocessor cores. These x86 microprocessor cores are needed to build ASICs, ASSPs, and SoCs for systems with graphical user interfaces and a connection to the Internet. No one is offering x86 cores. AMD, Intel, National, and STMicroelectronics are unlikely to sell cores that will compete with their chip offerings. That leaves Transmeta and Via Technologies. Both are fabless, so selling cores doesn't threaten fab capacity. It's an opportunity they shouldn't miss.

There's an emerging market for soft microprocessor cores for SoCs. These cores are needed to build ASICs, ASSPs, and SoCs for systems that don't have a graphical user interface. There's also an already-developed market for non-x86 microprocessor cores in applications with a graphical user interface. This is the market that ARM dominates. ARM has some years of opportunity as this market slowly shifts to x86.

The microprocessor market isn't monolithic; market segments can support a variety of Dynamic Silicon companies. Cypress Microsystems and Triscend, as they consolidate the microcontroller market, sell generic chips that are customized in the field to become SoCs. Altera and Xilinx build programmable logic devices with on-chip microprocessors. These chips are customized in the field to become SoCs. ARC, ARM, and Tensilica provide soft-core microprocessors that become the SoC's state sequencer.

A current joke says if you spent a thousand dollars on telecom stocks and a thousand dollars on beer last year then your empty beer cans are worth more than your telecom stock. Since sixty to seventy percent of sales for Altera and for Xilinx is in the "communications" sector, their stocks have been dragged down with the telecom stocks. Altera and Xilinx should do a better job of differentiating themselves from the telecom sector. Their businesses and their futures are more broadly based than they get credit for. Altera and Xilinx had a bad year last year, like everyone else. Revenues for Altera and for Xilinx dropped by almost forty percent last year. It sounds catastrophic, but it was the second largest year in Altera's history and the third largest for Xilinx. These companies sell the chips that are the future: generic chips with embedded microprocessor cores.

Nick Tredennick and Brion Shimamoto
July 16, 2002

NICK'S SCORECARD: WHO WINS, WHO LOSES

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Altera, Xilinx	Fabless	Excellent	Dominate the programmable logic business as the industry shifts toward generic chips that are customized in the field. Offering hard and soft microprocessor cores with programmable logic as the industry moves to SoC designs.
ARM Holdings, Ltd.	Fabless	Excellent	Dominates the soft core business. Has an early start in licensing and in royalty agreements for hard and soft cores. Fabless as the industry moves to fabless.
ARC International, Tensilica	Fabless	Good	Building soft cores as the industry moves to soft cores. Offering user-configurable resources and instructions as industry emphasis shifts to implementation efficiency. Fabless as the industry moves to fabless.
Cypress Microsystems, Triscend	Fabless	Good	Positioned to consolidate the microcontroller market. Building generic microcontrollers that are customized in the field as the industry moves to generic chip production and field customization.
National Semiconductor, STMicroelectronics		Integrated	Good Offering x86 microcontrollers as the market for them grows. Broad offerings in growing consumer electronics segments.
Via Technologies	Fabless	Good	Best position in x86 microprocessors for the "value" PC market. Needs to exploit opportunities in x86 microcontrollers and in x86 cores.
IBM	Integrated	OK	Offers cores, chips, design services, and fabrication. This integrated device manufacturer has begun to fragment its business.
MIPS	Fabless	OK	Dominates the market for set-top box cores. Will have to work to hold its market against invasion by x86 microcontrollers.
Transmeta	Fabless	OK	Great technology for the future, but needs to offer x86 microcontrollers and x86 cores in addition to its microprocessors. Positioned to benefit from growth in "value" PC market.
AMD	Integrated	Struggle	Offers high-end x86 microprocessors as the market shifts to mid-range and low-end microprocessors. Integrated device manufacturer as the industry fragments horizontally.
Intel	Integrated	Struggle	Offers high-end x86 microprocessors as the market shifts to mid-range and low-end microprocessors. Integrated device manufacturer as the industry fragments horizontally.
Motorola	Integrated	Struggle	Offers non-x86 microprocessors and microcontrollers as the market shifts from chips to cores. Integrated device manufacturer as the industry fragments horizontally.
Sun (Sparc)	Systems	Struggle	Its underperforming microprocessors and its systems occupy niches that are open to encroachment by x86-based systems.

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	6/28/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	13.60	12.70 - 33.60	5.2B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	29.70	26.60 - 52.74	10.9B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.27	£0.25 - £1.06	£115.6M
ARM Limited (ARMHY***)	Microprocessor and Systems-On-Chip Cores	11/26/01	16.59	6.52	5.81 - 19.20	2.2B
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	5.58	1.48 - 11.48	148.7M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	20.01	16.06 - 30.36	2.8B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	15.18	13.40 - 28.95	1.9B
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
Energy Conversion Devices (ENER)	Ovonic Unified Memory	6/18/02	27.69	15.69	12.64 - 28.00	343.6M
Foveon (none*)	CMOS Imaging Chips	6/18/02				
Microvision (MVIS)	MEMS-based Micro Displays, Nomad Head-Worn Display, Scanners	6/18/02	6.80	5.23	4.55 - 22.00	70.8M
National Semiconductor (NSM)	Geode x86 Microcontrollers, Consumer Orientation, 51% Ownership of Foveon	6/18/02	32.30	29.17	19.70 - 37.30	5.2B
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††	13.00	7.63 - 19.08	48.0B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	2.35	1.17 - 5.55	314.0M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	7.35	4.25 - 11.52	19.6B
VIA Technologies (2388.TW)	x86 Microprocessors for "Value" PCs	6/15/02	78.00		68.50 - 156.00	n/a
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	5.01	4.18 - 20.14	395.8M
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	22.43	19.52 - 47.16	7.5B

† 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.