GILDER TECHNOLOGY REPORT

February 2003 / Vol.IX No.2

Semiconductor Energy Crisis

As the world goes mobile, the microprocessor will lose its position as the workhorse in systems.

Inside:

The American auto industry chased raw horsepower, culminating in the muscle cars of the early 1970s, until the energy crisis changed the goal to efficiency. It was a surprise, and it was a traumatic change. The semiconductor industry is at a similar turning point. It's not that things are slowing and the problem is to decide where to turn off in a better direction. It's that the industry is going full speed down the interstate and sees "Road Closed" ahead. Talk about surprise! The industry has no brakes—who thought it needed any?

The industry has been motoring along doing one thing: making transistors smaller. I'm not talking about Moore's law; that's the *rate* at which small has been happening. I'm saying that *small no longer works*, because the rules have changed.

I'll recount enough history to show how lowering barriers by shrinking the transistor has become ingrained in the way semiconductor companies innovate.

Changes are sneaking up on us that challenge the microprocessor's role of workhorse in electronic systems.

Computers: an idea

The computer was a breakthrough *idea*—in problem solving. Before the computer, engineers solved problems by building special hardware. If hardware cost too much, they couldn't solve the problem.

Special hardware is expensive, efficient, and it lacks versatility. The computer is versatile, but it lacks efficiency. Programs (sequences of computer instructions) inefficiently translate algorithms tailored to solve a *particular* problem.

Fig. 1 is a microprocessor-based computer. The central processing unit (CPU) is a collection of expensive hardware. The relatively cheap memory stores problem solutions in the form of programs (sequences of instructions). The "bus" connecting the CPU and the memory carries instructions and data between them. Picture a huge number of commuters who live in sprawling-memory suburbs, who work in the tangled-CPU factory, and who suffer their commute over one narrow road.

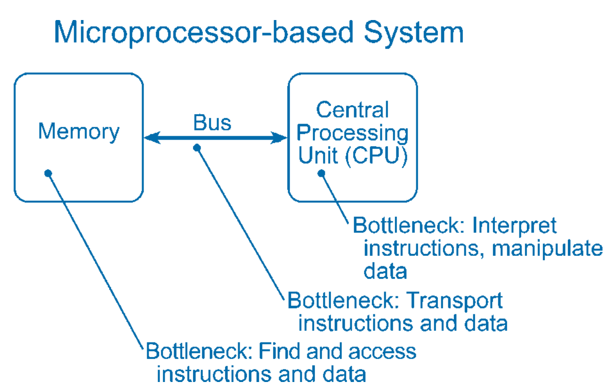

Fig. 1. The CPU, the memory, and the bus connecting them all have bottlenecks.

The first computers, called "mainframes," were big. These computers, from companies like **Burroughs**, **Control Data Corp.**, **Honeywell** (HON), **IBM** (IBM), **NCR** (NCR), and **Univac**, weighed tons and occupied entire rooms.

Integrated circuits (chips)

Integrated circuits made computers smaller, cheaper, faster, and more reliable. Integrated circuits lowered the barriers to building computer systems. (Building something that fits in a rack is easier than building something that needs a room.) "Minicomputers" entered the market. Companies like **Digital Equipment Corp.**, **Intergraph** (INGR), **Scientific Data Systems**, **Data General**, and **Prime** found a ready market. Mainframe manufacturers didn't see these new, barely capable computers as threats.

But minicomputers proliferated, transferring custody of data processing from corporate headquarters to departments.

Microprocessors

The computer was invented in the 1940s. The integrated circuit was invented in 1959. The year 1971 saw the two come together: the microprocessor was born. Engineers originally developed the microprocessor *to replace collections of integrated circuits.* The microprocessor became somewhat the ultimate integrated circuit, because it could *act* like any particular integrated circuit you needed. The microprocessor does this—similar to an actor taking cues—by fetching "cues"—instructions—from a memory. Each instruction *configures* the microprocessor to mimic, for an instant in time, some small hardware function. The microprocessor configures itself over and over, i.e., "runs" a sequence of instructions that mimics the behavior of any collection of discrete (separate) integrated circuits you would otherwise employ. Think of the microprocessor's instruction set as the collection of hardware functions it mimics. Instructions are like many small integrated circuits *timesharing* one large integrated circuit. Microprocessor-based implementations reduce component counts, and therefore cost, in systems where the performance of the microprocessor's high-speed mimicking is adequate.

By time-multiplexing their hardware, microprocessors achieve the effect of a full complement of hardware.

The shift from direct hardware implementations to microprocessor-based implementations enabled enormous growth in applications for two reasons. First, the microprocessor brought the computer's problem-solving method to embedded systems. It raised the level of abstraction in problem solving from hardware design to programming. This greatly increased the pool of designers and it made designers more productive, which accelerated the microprocessor's penetration into electronic systems. Second, the microprocessor consolidated the large number of integrated circuit types into a few major types (microprocessor, memory, input/output functions). Consolidating integrated-circuit types across applications led to high-volume production for the microprocessor; high volumes ensured both its declining cost and its rising performance.

The microprocessor, the icon of the industry's effort, has proved adequate for such a large number of applications that engineers, faced with a new problem or with escalating requirements, *begin* with microprocessor-based solutions. The microprocessor has been so successful that problem-solving methods based on it are entrenched in the engineering community, to the exclusion of other methods. Companies will, therefore, employ microprocessor-based solutions, whether or not they make sense. Universities teach microprocessor-based design almost exclusively. Microprocessor-based design has the momentum of a huge installed base of development systems and has the considerable backing of its many successful manufacturers.

But microprocessor-based systems have genetic bottlenecks inherited from the computer. To attempt to overcome the bottlenecks, manufacturers use high clock rates to speed the microprocessor's hardware; this burns lots of power. The microprocessor also forfeits efficiency because it is simulating hardware, imperfectly. (It can't be perfect—it has to have wide application.) These inefficiencies haven't mattered in the past. But now the microprocessor is becoming the sport-utility vehicle of integrated circuits: versatile and inefficient.

The battle for the desktop

The microprocessor's raison-d'être was to consolidate integrated circuits, but after ten years of Moore's-law advances the microprocessor became powerful enough to be the central processing unit in low-end computer systems. Continuing semiconductor progress transferred custody of data processing from departments to desktops. Minicomputer makers didn't see these new, barely capable desktop computers as threats.

In the beginning, personal-computer makers and workstation (non-x86-microprocessor-based) makers were vertically integrated. PC makers built for high-volume markets, that is, for consumers and for general business users. Workstation makers built for high-performance markets: science, engineering, and finance professionals.

Building for high-volume markets ensured that PC makers, as a group, could afford each Moore's-law turn-of-the-crank to boost PC performance and PC cost performance. Workstation makers built for performance-oriented markets and believed their performance would compel market share. Workstation makers didn't see these wimpy PCs as threats.

Intel's (INTC) volume-based strategy trounced the workstation makers' performance-based strategy. PC makers are forcing workstation makers into ever-smaller niches.

How did volume-oriented PCs overtake performance-oriented workstations in the workstations' own markets? The rapid performance improvement of the PC's microprocessor is part of the answer. The PC's microprocessor improved rapidly because of its high volume. There was always high demand at the PC's leading edge to encourage even higher performance designs. Also, Intel felt threatened by the workstation makers, so it fought to catch and to overtake their performance.

The larger cause, however, was that IBM lost control of the PC's design. Compaq and a host of other companies started to make "IBM PC compatibles." IBM's PC became the design standard; with the arrival of multiple sources, PC manufacturing began to fragment horizontally. **Phoenix** (PTEC), **Microsoft** (MSFT) and others reproduced the basic input/output system (BIOS), wrote operating systems, and wrote applications. Countless manufacturers made hard disks, CD-ROMs, floppy drives, power supplies, video cards, displays, keyboards, and printers. PC makers turned into assemblers—sourcing standard components to build systems. Meanwhile, workstation makers remained vertically integrated. They built their own microprocessors, designed their own memory systems and video cards, and they wrote their own BIOS and their own variants of the Unix operating system. Each workstation manufacturer was a tower of vertical development competing for a small workstation market. Meanwhile, hardware and software development in the huge PC market was shared across system makers. Result: PCs were much cheaper and they evolved much faster.

Workstation makers amortized research and development costs across tens of thousands of units. PC makers amortized costs across tens of millions of units. In those days, leading-edge microprocessor development might cost fifty million dollars. Amortized development cost added a dollar or so to the PC-microprocessor's cost. Amortized development cost added a thousand dollars or so to the workstation-microprocessor's cost. This workstation premium, based on lower volumes, extended to all the system's hardware and software components—a huge pricing advantage for PC makers.

It's more than fifty years since the first commercial computer, more than thirty years since the microprocessor was invented, and more than twenty since its use in workstations and PCs. We've come to believe in instruction-based processing—programming—as *the way* to solve problems. We've forgotten the microprocessor's origin. We're about to be reminded. The idea of the computer in the form of the microprocessor has succeeded because its inefficiencies haven't mattered. That's changing.

Tethered and untethered, powerful devices

The world is splitting into tethered and untethered systems. Tethered systems connect to wall outlets for power and make up the world's physically interconnected collection of computing, access ports, data transport networks, and storage. Untethered systems use onboard power and are both collectors and consumers of data. Untethered systems are on the front lines of interaction with the physical world. The PC is primarily a tethered system. As a consumer item, it balances cost with performance. Emerging untethered systems add power conservation to the mix. *This changes the design objective from cost performance to cost-performance-per-watt.*

Semiconductor progress is shrinking the computer again, this time from desktop/laptop to shirt pocket. PC makers won't see these barely capable shirt-pocket computers as threats. And this time they won't be. The desktop computer is here to stay. The PC matured when its performance exceeded the demands of most users. Once that happened, the market shifted from leading-edge PCs to "value PCs."

The change in the market from high-margin, high-performance PCs to lower-margin, high-value PCs will shift engineering emphasis to the design of (more profitable) untethered systems. The problem is that microprocessor-based solutions aren't efficient enough to satisfy the cost-performance-per-watt requirements of untethered systems. The market needs the versatility of microprocessors (suits many roles) with the efficiency of ASICs (special hardware).

Moving computing to the shirt pocket (e.g., cell phones, personal-digital assistants, GPS receivers, digital cameras, MP3 players, game players) changes two system characteristics. First, the design goal becomes cost-performance-per-watt. The second change, subtle but perhaps more important, is that these devices, as collectors and consumers of data, *move humans to the end of the process.* For the history of computing, humans have stood *between* the data and the computing: collecting data and feeding it to the computer and collecting and interpreting results. Humans have been the computer's interface to the physical world. With shirt-pocket computers, direct links are being established between the computer and the physical world. The human moves to a position as observer and manager.

Microprocessors are no longer the answer

As the microprocessor got faster, its range of applications expanded. But, as fig. 1 shows, a microprocessor-based system is a collection of bottlenecks. The microprocessor is bottlenecked because its instructions timeshare its hardware. (The microprocessor is like a hotel where guests (instructions) show up. The overhead to prepare for and to clean up after guests greatly exceeds what has to be done during their stay.) The interface between the microprocessor and its memory is bottlenecked by sharing the transport path for instructions. (Traffic clogs the one road to the hotel.) The memory is bottlenecked for finding and accessing instructions and data, because the memory is good at holding things, but is slow (keeps the microprocessor waiting) at finding them.

The shift from cost performance to cost-performance-per-watt may not seem like a big deal. Instead of designing for maximum performance, microprocessor makers can design in a way that balances power use and performance. That's what the microprocessor makers are doing. But remember where the microprocessor came from—it has always been an inefficient substitute for direct hardware implementation in applications where the microprocessor had adequate performance. It timeshared its hardware to reduce overall cost. Manufacturers increased the microprocessor's performance by speeding the timesharing (more customers, shorter stays, in the hotel analogy) of its hardware resources. Speeding its timesharing does nothing for *efficiency*. Emerging untethered systems want efficiency. Microprocessor-based implementations can't get there. But the business models and company cultures of the leading microprocessor companies confine them to instruction-based solutions.

TELECOSM TECHNOLOGIES

Ciena	(CIEN)
Corvis	(CORV)
JDS Uniphase	(JDSU)
Avanex	(AVNX)
Essex	(ESEX.OB)
Equinix	(EQIX)
Sprint PCS	(PCS)
Qualcomm	(QCOM)
Broadcom	(BRCM)
Altera	(ALTR)
EZchip	(LNOP)
Terayon	(TERN)
National Semiconductor	(NSM)
Synaptics	(SYNA)
Intel	(INTC)
Soma Networks	(PRIVATE)
Narad Networks	(PRIVATE)
Flextronics	(FLEX)
Taiwan Semiconductor	(TSM)
Transmeta	(TMTA)
Analog Devices	(ADI)
ARC Cores	(ARK)
ARM Limited	(ARMHY)
Calient	(PRIVATE)
Celoxica	(PRIVATE)
Cepheid	(CPHD)
Chartered Semiconductor	(CHRT)
Coventor	(PRIVATE)
Cypress	(CY)
Cyrano Sciences	(PRIVATE)
Energy Conversion Devices	(ENER)
Foveon	(PRIVATE)
Legend Group Limited	(LGHLY.PK)
Microvision	(MVIS)
QuickSilver Technology	(PRIVATE)
SiRF	(PRIVATE)
Tensilica	(PRIVATE)
Triscend	(Private)
United Microelectronics	(UMC)
VIA Technologies	(2388.TW)
Wind River Systems	(WIND)
Xilinx	(XLNX)

Corvis (CORV)

WDM SYSTEMS, RAMAN AMPLIFICATION, EDGE SWITCHES

FEBRUARY 19: 0.65 52-WEEK RANGE 0.47-1.86 MARKET CAP: 268M

Revenues of $7.1 million in the December quarter are a clear improvement from last quarter's $1.4 million in sales. More vital for the company at this time is its continued focus on reducing costs and hence protecting its strong balance sheet. Corvis ended the quarter with $504 million in cash, with cash burn coming in at $44 million. Despite the fact that the majority of that cash use is related to success-based carrier trials, that number needs to be reduced. Reducing cash burn to the company's target of $25 million per quarter would allow the current cash to last through 2006. Such strong focus on cash preservation is a necessity due to the continued uncertainty surrounding the return of carrier capital expenditures. However, we do see some signs of life. First, Tier 1 carrier trial activity remains robust. AT&T has experienced a clear flight to quality, in terms of network traffic growth, and it is believed that they will need to deploy new next-generation gear in 2003 to cope with capacity restraints. Corvis is seen as one of two realistic candidates, the other being Ciena. Second, Corvis has successfully completed its first undersea trial. A subsequent contract win would most likely be for capacity (wavelength) additions to a large incumbent network. Third, Corvis announced that it shipped product to the U.S. government in Q402, with revenue recognition to follow in Q103. Corvis is well positioned to win the ultra long-haul portion of the GIG-BE government project which is estimated to be worth $200-$400 million over two to three years. These signs of hope, coupled with the fact that Corvis now trades at nearly a 50% discount to cash, represent an opportunity to the high-risk, patient investor.

JDS Uniphase (JDSU)

ACTIVE AND PASSIVE OPTICAL COMPONENTS

FEBRUARY 19: 3.00 52-WEEK RANGE 1.58-6.90 MARKET CAP: 4.1B

GOOD NEWS, BAD NEWS—CIBC has essentially called the bottom on JDSU's business prospects and has outlined several items that will drive a recovery. These include: the potential for a sequential sales increase in the March quarter; a telecom capital spending "budget flush" in Q4, with some spending spilling over into Q1; an AT&T ramp in long-haul spending in 2H03; two sizable Chinese long-haul builds; OEM customers like LU and NT sidestepping liquidity issues; long-term investors seeking a strong balance sheet; and the basic oversold situation in October. The bad news, however, is that CIBC feels these highlights may already be priced into the stock.

Avanex (AVNX)

ADAPTIVE PHOTONIC PROCESSORS

FEBRUARY 19: 0.83 52-WEEK RANGE 0.63-5.20 MARKET CAP: 56M

Avanex reported revenues of $5.3 million, up from $5.2 million last quarter, as the company continues to focus on cost cutting through the consolidation of its operations in Fremont, California. Customers in the quarter increased to 16 from 12 in the previous quarter. Metro DWDM and OADM applications continue to drive revenues, accounting for more than 70% of total revenues this quarter. The company recognized revenues from newer products and technologies, such as the PowerShaper FDS chromatic dispersion compensator. While hurting gross margins, the superior performance and low cost of these new products will drive the company's future. Cash burn in the quarter was $11 million, putting total cash now at $151 million. Cash burn estimates have the company exiting FY04 (May 04) with ample net cash of about $105 million, or $1.50 per share.

StorageNetworks (STOR)

DATA STORAGE MANAGEMENT, SOFTWARE

FEBRUARY 19: 0.92 52-WEEK RANGE 0.78-4.75 MARKET CAP: 89M

StorageNetworks entered the Gilder Paradigm as a remote storage management solutions provider, morphed into a traditional storage services provider (SSP), and now, following the January 30 release of its earnings and subsequent conference call, leaves us guessing as to where they plan to venture next. STOR's major asset throughout the downturn has been its strong cash position, around $186 million. That bastion is now endangered by the company's announcement that it intends to acquire another business in order to move forward. These uncertainties lead us to remove StorageNetworks from the "list" of Telecosm Technologies.

Sprint PCS (PCS)

NATIONWIDE CDMA WIRELESS NETWORK

FEBRUARY 19: 4.19 52-WEEK RANGE 1.75-13.45 MARKET CAP: 4.3B

While everyone expected the worst, Sprint PCS delivered a nice upside surprise with a return to positive net subscriber additions of 250,000 in the quarter. One focus of the company conference call was the continued benefits of the CDMA2000 1x overlay, in the form of improved network efficiencies, as the 1x handset base approached 50% of total subscribers. Vision data customers now number 630,000, and during the quarter approximately 50% of the data subs were new to PCS, while 50% upgraded their service. The company currently has two-thirds of the Vision data customers continuing with the service after the three-month free trial is over. Additional areas of management concentration, especially for PCS president Len Lauer, included continued improvement in customer quality (prime) and in customer service. Greater than 60% of this quarter's adds were prime customers, bringing the Sprint PCS subscriber base to 73% prime. Customer complaints dropped 15% q/q, and the percentage of satisfied customers increased to 71% from 64%.

Qualcomm (QCOM)

CDMA INTEGRATED CIRCUITS, IP, SOFTWARE

FEBRUARY 19: 34.70 52-WEEK RANGE 23.21-44.65 MARKET CAP: 28B

RESULTS—Qualcomm simply continues to deliver tremendous results in a far from simple environment, shipping 29 million chipsets this quarter, up from 20 million last quarter. Of those 29 million chipsets, over 22 million were CDMA2000 1x, while 700,000 were CDMA2000 1x EV-DO. As of December 2002, South Korea counted 175,000 subscribers on its two EV-DO networks owned by SKT and KTF. For 2003, the South Korean operators are targeting 5.2 million EV-DO subscribers. Uptake is expected to be driven by dropping prices for EV-DO handsets and through advanced new features such as camcorder phones, high-resolution cameras, and MPEG4 chips. KTF is about to launch a ten-channel TV broadcasting service, while SKT is already offering made-for-handset short movies. Wireless LAN and EV-DO bundling is also expected shortly from KTF. Japan's KDDI, hot off the success of its 1x network, plans to launch 1x EV-DO in October.

Rumors: China Telecom, the Chinese operator that is supposedly launching commercial 3G services using TD-SCDMA or WCDMA, depending on who you ask, is testing another form of 3G—CDMA2000—at 450 MHZ, confirming rumors that have been circulating for months.

Texas Instruments (TXN)

DIGITAL, ANALOG, MIXED-SIGNAL PROCESSORS

FEBRUARY 19: 16.24 52-WEEK RANGE 13.10-35.94 MARKET CAP: 28B

Texas Instruments is the market leader in DSPs (digital signal processors). TI is also an integrated device manufacturer (IDM), which means it is vertically integrated for chip design, manufacturing, and sales. DSPs use the computer's instruction-based model for signal processing. Instruction-based processing is inherently inefficient. It simulates functions rather than implementing them directly. That's OK if the processing task isn't demanding and it's OK if power is free. Instruction-based processing also perpetuates the processor-memory bottleneck. Engineers strive to increase the DSP's performance while their memory-chip counterparts strive to increase memory capacity. This difference in emphasis (speed vs. capacity) results in a growing mismatch of processor and memory speed. Moreover, power is anything but free for mobile applications. DSPs and microprocessors have historically traded performance for voltage, lowering the supply voltage to control power as frequency rises. Voltages are now below one volt, so there's little room left to trade voltage for higher clock rates. In mobile applications, therefore, direct implementation of functions will win. DSPs have neither the efficiency nor the performance to meet future requirements. As TI holds to its integrated business model as the rest of the industry fragments horizontally, the company will struggle. For all but a few high-volume players such as Intel, the IDM business model will prove too expensive to maintain. But long-time IDMs like TI have an entrenched culture that makes transition to horizontal businesses difficult. None of this will happen tomorrow or even next month. TI is huge, with an enormous installed base of development systems used to design-in its DSPs. Microprocessor- and DSP-based design methods are entrenched in the engineering community. University programs continue to teach instruction-based design methods. But the fact is that signal processing has a bright future, but DSPs don't. For these reasons as well as the argument in this month's letter, we have removed Texas Instruments from the "list" of Telecosm Technologies.

National Semiconductor (NSM)

SINGLE-CHIP SYSTEMS, FOVEON IMAGERS

FEBRUARY 19: 14.91 52-WEEK RANGE 9.95-37.30 MARKET CAP: 2.6B

A key component of National's growth story is increasing semiconductor content in all of the vertical market segments that the company addresses. Displays continue to be the best performing end-market for NSM. Of the eleven different chips used in flat panel displays (FPD), National provides eight. While NSM supplies $1-$2 of content in an older CRT, the company supplies an average of $7 in content for the average FPD and up to $20-$35 in a high-end design.

Soma Networks

BROADBAND WIRELESS ACCESS, NETWORK SOFTWARE

PRIVATE

Soma has continued to fly under the radar, dodging the spotlight sought and attained by many of its competitors in the broadband wireless space. However, Soma is delivering the signal while companies such as Flarion and Navini continue providing the noise. Bank of America awarded Soma's progress with a $10 million financing agreement which will guide the company as it enters the large-scale deployment phase of its operations.

Transmeta (TMTA)

MICROPROCESSOR INSTRUCTION SETS

FEBRUARY 19: 1.35 52-WEEK RANGE 0.74-4.47 MARKET CAP: 179M

Two recent articles in DigiTimes, coupled with earlier announcements of design wins, solidify that Transmeta is gaining traction and moving in the right direction. It was reported that Transmeta is to become the dominant supplier for desktop replacement notebooks from Elitegroup Computer Systems (ECS), a large Taiwan-based manufacturer. This development would push VIA Technologies out of the top spot at ECS in the desknote market and eliminate Asian desknote sales of Advanced Micro Devices at ECS altogether. According to the report, ECS would use Transmeta for 60% of its desknotes, or roughly 360,000 units. DigiTimes also reported that Transmeta has received a follow-on order from Hewlett-Packard for its latest Compaq Tablet PC, expected to launch at the end of Q103.

Flextronics (FLEX)

CONTRACT MANUFACTURING

FEBRUARY 19: 8.28 52-WEEK RANGE 5.47-20.47 MARKET CAP: 4.2B

Flextronics reported a better than expected December quarter with revenues of $3.85 billion. The mobile phone outsourcing trend we highlighted last month was a major contributor: the Sony-Ericsson relationship alone represented 12% of total sales, up 40% sequentially. Analysts and investors have been concerned over Ericsson's commitment to the joint venture. Therefore, the announcement by Sony and Ericsson that the two OEMs had decided to each contribute 150 million euros to the joint venture helped lessen fears. The company's ODM initiatives received more discussion during the quarter as examples of greater standardization from reference designs, integrated chipsets, and operating systems that allow for commoditization of the handset, similar to the PC. Last year, Flextronics introduced PhoneOne, its ODM-related mobile phone, which is targeted at the lower-end of the handset market. The company believes that it can reduce the cost of making a similar mobile handset by 25%. Flextronics announced a co-development agreement for its PhoneOne with Qualcomm, clearly aimed at the developing markets of China and India. Qualcomm has squarely pinned its fortunes on these regions in 2003, and Flextronics could be a key enabler and beneficiary.

Taiwan Semiconductor (TSM)

CMOS SEMICONDUCTOR FOUNDRY

FEBRUARY 19: 6.85 52-WEEK RANGE 5.31-19.08 MARKET CAP: 25.8B

Taiwan Semiconductor Manufacturing Corporation's January sales rebounded dramatically, up 16% month-over-month, after December sales declined sharply, falling 22%. The company managed to grow revenues 28% year-over-year, while the overall semiconductor industry remained flat. Fourth quarter results provided further evidence of the slowing adoption rate for next-generation semiconductor processes. During the most recent quarter, 250 nm products represented the lion's share of revenues at 25%, while 130 nm or less equaled 8%.

Why microprocessors won't do

Two things are sneaking up on the industry.

First, the PC, which has led the semiconductor industry for a long while, is about to cede its position. The value PC is displacing the leading-edge PC as the bulk of the market. Engineering resources once dedicated to increasing the PC's performance will be reallocated to more profitable projects.

Second is the rise of untethered systems, which bring with them the requirement to optimize cost-performance-per-watt. Optimizing cost-performance-per-watt requires circuit efficiency. That's a mismatch for the microprocessor, which is based on power-inefficient timesharing and on raising the level of abstraction in problem solving through inherently inefficient instruction-based processing. The microprocessor has been raising its clock speed to boost performance. But doubling clock speed doubles power use. Designers compensate by lowering the microprocessor's supply voltage. Cutting the supply voltage in half lets the microprocessor run four times faster for the same power. That works until designers run out of room to lower voltage. We're there, there's little room left. But speeding clocks and lowering voltage do nothing to fix the microprocessor's inherent inefficiencies.

If microprocessor-based solutions won't do and the companies that make them are too rigid to explore alternatives, then there's opportunity. But for whom? Application-specific integrated circuits (ASICs) come to mind: they are direct hardware implementations, so they can have the necessary efficiency and performance. In high-volume applications, perhaps they could meet the cost requirements. But ASICs are too inflexible to meet the evolving needs of emerging markets.

The long answer

The answer is difficult because the semiconductor industry is at a turning point. The industry cannot just improve what it already knows how to do (microprocessors, digital signal processors, Moore's law); it has to do something *new* (power-efficient design). And the industry's response is "I don't know." We will wait while the industry postures and scrambles.

The microprocessor's track record of burgeoning performance has ingrained the raising of clock speed as *the way* to advance. There's an analogous story for the semiconductor manufacturing process. Moving to ever-smaller transistors is a forty-year success recipe. Building smaller transistors at each generation is established as *the way* to advance semiconductor process. *Over time, the path to success becomes part of the company's culture and a part of its customers' expectations.*

Manufacturers build microprocessors with faster and faster clocks because that has worked for thirty years. The microprocessor's customers expect faster and faster clocks. Chip makers build smaller and smaller transistors because that has worked for forty years. Building smaller transistors is built into the business models of the integrated device manufacturers. We feel the effect of ingrained customer expectations in the personal computer business. We buy PCs with faster microprocessor clocks (that only marginally improve *system* performance) because we established our buying habits when higher clock speeds *did* improve system performance.

Culture and expectations change slowly; behavior persists beyond the events that should precipitate change. That's where we are today with the microprocessor, with the semiconductor process, and with programmable logic devices. The events that precipitate change have occurred, but, because today's methods are ingrained in company culture and in customer expectations, the transition to new methods and to new products will be slow.

With Intel as the dominant company making the dominant microprocessor for the dominant (PC) platform, Intel's company culture greatly affects the semiconductor industry. "Hive" and "one-minded" are words that describe Intel's culture. In the industry, Intel is known as the world's largest single-celled organism. Intel's bumper sticker reads "Moore's law and microprocessors forever." The word "alternative" is not in Intel's vocabulary.

Contrast this with IBM, a company with a contention culture. (Outsiders never see this because marketing communications is an IBM core competence.) Every IBM product direction is surrounded by technical alternatives. Proponents *pray* that their alternative will sink the incumbent. The atmosphere is reminiscent of courtrooms. This is *normal* IBM culture. While most credit Lou Gerstner, IBM's contention culture played a major role in saving the company, enabling it to change direction.

PLDs today

Perhaps programmable logic holds the answer. The programmable logic device (PLD) is conceptually a two-layer device. One layer is configurable wiring and configurable logic blocks. The second layer is the configuration memory. Bits in the configuration memory personalize each logic block and specify how the wiring connects the logic blocks to build a direct hardware implementation of larger functions. This arrangement seems to have the versatility of the microprocessor with the efficiency of the ASIC's direct hardware implementation. But today's PLDs are too big, they're too slow, they're too expensive, and they have more overhead than the U.S. government.

I said programmable logic might be the answer and followed that with a crippling list of problems. Let's square that. The PLD's promise lies in its vision; its problems lie in its current implementations. The PLD's vision includes the efficiency of direct hardware implementation and the flexibility of rapid reconfiguration (changing from one direct hardware implementation to another).

We are at the same place with PLDs that we were with the *computer* in the 1950s. People only saw room-size computers, not the vision of computing.

The difficulties in applying PLDs' current products to new applications lie in the current manufacturers' cultural legacies and in their customers' expectations. **Altera** (ALTR) and **Xilinx** (XLNX) dominate today's PLD markets. PLD makers did the same thing that the microprocessor companies and semiconductor manufacturers did; they built a path to success on a model that worked. PLD makers began by making programmable logic chips that: 1) consolidated the "glue logic" in the design of circuit boards and 2) could be used for prototyping (building the first hardware of a circuit). Glue logic is miscellaneous logic (e.g., invert the current signal, send eight copies to other parts of the circuit, combine three signals from other circuits) that connects major components in a system.

Customers' needs determined the capabilities of today's PLDs. Prototyping chips didn't need fast configuration and they didn't have to be small and fast or even cheap, since they wouldn't be the final parts that went into production systems.

The shift to programmable logic

I doubt that the shift from microprocessor-based implementations to programmable logic implementations will enjoy the same ease of market penetration or the same rapid growth as the shift from discrete integrated-circuit implementations to microprocessor-based implementations. First, the microprocessor easily invaded hardware systems because it complemented still-present hardware. Incumbent hardware didn't fight the microprocessor's advance.

Microprocessor makers don't yet see programmable logic as a threat. As in the past, the newcomers' problems are big enough to let them fly under incumbents' radar. But the microprocessor occupies applications that programmable logic is invading. The microprocessor *will* find itself competing with programmable logic. It's a losing battle for the microprocessor, but it will slow programmable logic's advance. Second, the microprocessor boosted the number of designers and it raised their productivity. Programmable logic doesn't have that advantage. In fact, programmable-logic-based design requires different skills than microprocessor-based design.

Meanwhile, the world will soon be awash in cheap electrical engineering graduates. Countries such as China and India graduate ten times as many electrical engineers as the United States. This is an opportunity for the non-U.S. tech industry to score a coup.

Programmable logic needs a few proof-of-concept implementations to prove its worth in untethered applications. Once that happens, it will boost financial backing for development software from companies such as **AccelChip**, **Celoxica**, and **MathWorks**. More mature development software will raise the level of abstraction and enable more designers. Like the microprocessor, programmable logic is generic in manufacture and customized in the field. High-volume manufacturing will ensure its declining cost and its rising performance. Its advantage over the microprocessor is the efficiency of direct hardware implementation.

The microprocessor was invented for its role in embedded applications. Today's programmable logic devices weren't invented for the job they're about to take on in untethered systems. Today's PLDs aren't suited to the application. But the big PLD makers see the opportunity for growth in a market that's *ten times the size of their current market*, so they are changing their components to meet the needs of these new markets. Their first objectives are microprocessors and digital signal processors in tethered applications, where current PLDs already meet the needs of many applications.

The development of programmable logic chips that are customized for untethered applications has been left to start-up companies such as **Ascenium**, **QuickSilver Technology, Inc.**, and **Streamatics**. Even though it's a difficult time to get venture funding, the companies' timing is good for the opportunity—not too early, not too late. The market for untethered applications is emerging and should grow rapidly. At the same time, incumbent microprocessors and digital signal processors are unsuited to the task. Microprocessors and digital signal processors have traded all the voltage they can for additional performance, so they won't have enough absolute computing power to meet the needs of emerging applications, such as advanced cell phone protocols. Programmable logic, because it is direct hardware implementation, has the performance for emerging applications. Also, microprocessors and digital signal processors aren't power-efficient enough to compete with programmable logic implementations. Programmable logic is emerging as the incumbent chips reach their limits.

Manufacturers ship billions of microprocessors a year. The microprocessor and its makers won't disappear tomorrow. In fact, the market for microprocessors will grow. The microprocessor won't even be kicked out of untethered systems. The microprocessor will only lose its job as the *workhorse* in untethered systems. It retains its vital function as manager—in deciding what happens when.

Whence the new innovation model?

Integrated circuits lowered the price of entry into the computer business by making computers tractable (not room-size).

- The microprocessor brought programmers to problem solving on a massive scale.
- The PC was the first horizontally integrated computer *industry*.

The stars of the semiconductor story—the integrated circuit, the microprocessor, and the personal computer—have sustained two themes. The first theme has been Moore's law for making semiconductors. The second theme has been the PC as *the* platform. *The semiconductor industry's innovation* *model has been to lower barriers to entry, embracing these themes.* We have reached (in the case of the value PC) or are reaching (in the case of the value transistor) the end. Microprocessors with more gigahertz and transistors with smaller line widths are no longer in the catbird seat. The appearance of the value transistor (*Dynamic Silicon*, December 2002) and the market shift to untethered systems is our wake-up call to deemphasize our current cast of stars.

Since the invention of the integrated circuit, the transistor hasn't been good enough. Moore's-law progress has been making transistors better—and in so doing, smaller transistors lowered barriers that protected businesses such as mainframes, minicomputers, and workstations. Business as usual in semiconductor companies is the turmoil caused by lowering barriers. That's worked for *four decades* of electronic systems. The semiconductor industry has built enormous momentum following this model. But Moore's-law progress in microprocessors has made the personal computer good enough, and Moore's-law progress in transistors has built transistors that are good enough. The value PC and the value transistor change the rules. No more business-as-usual lowering barriers: the value PC is the limit of tethered systems. Computing is moving to the shirt pocket. Untethered systems want cost performance, but they also want power efficiency. The microprocessor can't do it. Perhaps programmable logic can, but the whole industry is lumbering forward on its own inertia, like a supertanker that didn't expect to have to turn.

Energy Crisis Scorecard: Who Wins, Who Loses

COMPANY	TYPE OF COMPANY	FUTURE POSITION	THE WAY I SEE IT
Altera, Xilinx	Fabless	Excellent	Altera and Xilinx dominate the growing market for programmable logic chips. Their immediate opportunity is displacing microprocessors and digital signal processors in performance-oriented tethered systems.
TSMC, UMC	Foundry	Excellent	TSMC and UMC are the primary chip manufacterers for the leading programmable logic vendors. The growing market for programmable logic chips helps offset rising equipment, process development, and mask costs with larger production runs of general-purpose chips.
AccelChip, Celoxica, MathWorks, Synopsys	Development Software	Good	Design with programmable logic offers efficiency and performance needed for untethered systems, but is more difficult than microprocessor-based design. Software that maps algorithms into programmable logic will gain market share with the rise of untethered systems.
ARC International, Tensilica	Fabless	Good	Configurable microprocessors bridge the gap between fixed-instruction-set micro processors and programmable logic. Configurable microprocessors offer huge performance advantages over their fixed-instruction relatives while retaining familiar microprocessor-based design methods.
ARM	Fabless	Good	ARM's soft-core microprocessors, which are entrenched in untethered applications, will benefit from growth in the market for shirt-pocket computing.
Ascenium, GateChange, QuickSilver	Fabless	Good	There is opportunity designing programmable logic devices suited to untethered applications, but displacing entrenched microprocessors and digital signal processors will be difficult.
MemoryLogix, Transmeta, VIA Technologies	Fabless	OK	There is opportunity for soft-core x86 in the emerging market for untethered systems. These companies could offer soft-core x86 microprocessors, though none currently does so.
Intel, Motorola, Texas Instruments	Integrated Device Manufacturer	Struggle	Integrated device manufacturers continue to build performance-oriented micro-processors and digital signal processors as the market shifts from cost performance to cost-performance-per-watt. They will continue to build instruction-based processors as the market shifts to programmable logic.
Sun Microsystems	Computer Systems	Fail	Vertically integrated computer makers pay all the costs of hardware and software development but occupy only a small market position in workstations and servers. They cannot compete with the horizontally fragmented x86-computer business that amortizes costs across substantially higher volumes.

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.

GILDER TECHNOLOGY REPORT PUBLISHED BY GILDER PUBLISHING, LLC AND FORBES INC. • Copyright ©2003, by Gilder Publishing, LLC

291A MAIN STREET, GREAT BARRINGTON, MA 01230, TEL: (413) 644-2100, FAX: (413) 644-2123 • EMAIL: INFO@GILDERTECH.COM

EDITOR IN CHIEF George Gilder
EDITORS Nick Tredennick, Brion Shimamoto
EXECUTIVE EDITOR Bret Swanson
ANALYSTS Mary Gorski, John Hammill
MANAGING EDITOR Marie Lavinio
ART DIRECTOR Charles Bork
SUBSCRIPTION DIRECTOR Rosaline Fernandes

For subscription information telephone toll free: (800) 292-4380 • WWW.GILDERTECH.COM