GILDER TECHNOLOGY REPORT

August 2003 / Vol. IX No. 8

From Programmed To Reconfigurable

The first generic semiconductor was the transistor, the second was the microprocessor, soon it will be the PLD's turn.

Inside:

The PLD (programmable logic device) market doesn't draw much attention. It's too small. As the dominant companies saturate the PLD market, they look to the microprocessor market and to the ASIC (application-specific integrated circuit) and ASSP (application-specific standard product) market—each at least ten times the size of the PLD market—for new opportunities. This is the story of how they will shift their strategy to invade these new markets. And it's the story of the market forces that will aid them, the market forces that will oppose and inhibit them, and of how they will fare.

Semiconductor market

We think of the personal computer in terms of its microprocessor and its memory. But look at a PC system board sometime; it's dominated by other stuff—so much stuff that it's difficult to find the microprocessor and memory. The microprocessor is the chip with a Coke-can-size heat sink and fan on top of it and the memory is a few chewing-gum-size boards mounted on edge nearby. What's all the other stuff? It's analog and digital logic and it's discrete electrical components (resistors, capacitors, inductors, and switches).

The semiconductor market for 2003 will be $175 billion. Microprocessors will be $47 B and digital logic will be $38 B. The microprocessor segment includes digital signal processors (DSPs). The balance is memory, analog, discrete, and optical components. Thirty-five billion dollars of the digital logic is application-specific integrated circuits and application-specific standard products. Less than $3 B is programmable logic devices.

I'm a programmable logic enthusiast. What is programmable logic? And why should anyone care about a segment that's 5 percent of the semiconductor market?

PLD applications fall into two categories. At the low end, PLDs consolidate the system board's miscellaneous digital logic, called "glue logic," into a single chip or into a few chips. Mopping up miscellaneous digital logic reduces cost and improves reliability. At the high end, engineers prototype complex circuits with PLDs before the logic is committed to expensive ASICs or ASSPs. PLD-based prototypes are much closer to the final implementation than a computer-simulated circuit and they run hundreds or thousands of times faster. At the same time, they are cheaper and more easily modified and debugged than an ASIC or ASSP.

PLDs consist of general-purpose logic, wires to connect the logic, and a personalization memory. Bits in the personalization memory control the connections between wires and logic to build circuits. PLDs are generic when manufactured and they are customized in the field (by loading bits in the personalization memory). This sounds quaint, but, as we shall see, it is crucial to the PLD's future success.

CONTINUED ON PAGE 2

Altera (ALTR) and **Xilinx** (XLNX) dominate the market for programmable logic devices. Altera is celebrating its twentieth anniversary in 2003; Xilinx will celebrate its twentieth next year. These companies emerged as the market for glue logic was maturing and as the market for embedded microprocessors was growing rapidly. They built their businesses initially on consolidating glue logic into general-purpose programmable chips that sat next to the microprocessor in an embedded system. They are strong companies with a twenty-year history of rapid growth, no debt, and plenty of cash. Margins on their leading-edge chips are as high as Intel's margins on its high-end microprocessors.

ASICs and ASSPs

An application-specific integrated circuit is a circuit-based chip for a particular application. Cell phones and digital televisions are two examples. ASICs are fast and efficient, but the speed and efficiency come at a cost. The masks (patterns for building circuits on a chip) are expensive. The mask set for a 130-nm process can cost $600,000. And costs rise rapidly with each semiconductor process generation. The mask set for a 90-nm process can cost $1,300,000. Engineering costs (designing the circuits that the mask set represents) can be ten times the mask cost. These are *fixed costs*; they are the same whether the market is one chip or a hundred million. Because of high fixed costs, ASICs are appropriate for high-volume applications. Leading ASIC vendors include **Agere Systems** (AGRb), **Agilent** (A), **Fujitsu** (FJTSY.PK), **IBM** (IBM), **LSI Logic** (LSI), **NEC** (NIPNY), **STMicroelectronics** (STM), and **Toshiba** (TOSBF.PK).

An application-specific standard product is a chip built for a single application. That sounds just like an ASIC, and it is except that an ASIC is built for one application by a single manufacturer. **Nokia** (NOK), for example, builds ASICs for its cell phones and does not sell them to competitors. **Analog Devices** (ADI) and **Texas Instruments** (TXN), on the other hand, build ASSPs for cell phones, and will sell them to anyone. ASIC fixed costs amortize across the unit volume of a single manufacturer. ASSP fixed costs amortize across the unit volumes of many manufacturers. The ASSP's advantage is lower cost, and its disadvantage is that competitors have access to the same chip, making product differentiation more difficult.

Problems with ASICs

Fixed costs (mask cost and engineering cost) approximately double from one process generation to the next, so the market for the end product has to double in units for the amortized fixed cost to remain constant. Conclusion: more advanced processes require larger markets.

The leading-edge of performance and of transistor capacity that an ASIC can supply doubles every eighteen months (Moore's law). The demand for performance and for transistor capacity spreads (across types of users) with time and this demand grows at a rate that is slower than Moore's law. That creates a problem for ASICs. Once upon a time, ASIC makers didn't worry about encroachment by PLD makers because ASICs were ten times faster and had more than ten times the number of transistors. PLD performance and transistor capacity didn't pose a threat. But PLDs have been improving at the Moore's-law rate. Since the *demand* for performance and for transistor capacity spreads and grows at a slower rate, the supply curve eventually pokes into the demand region (see fig. 1). Eventually is now. PLDs are encroaching on application territory that was once the exclusive property of ASICs. For the majority of applications, ASIC performance and transistor capacity are improving *beyond* demand.

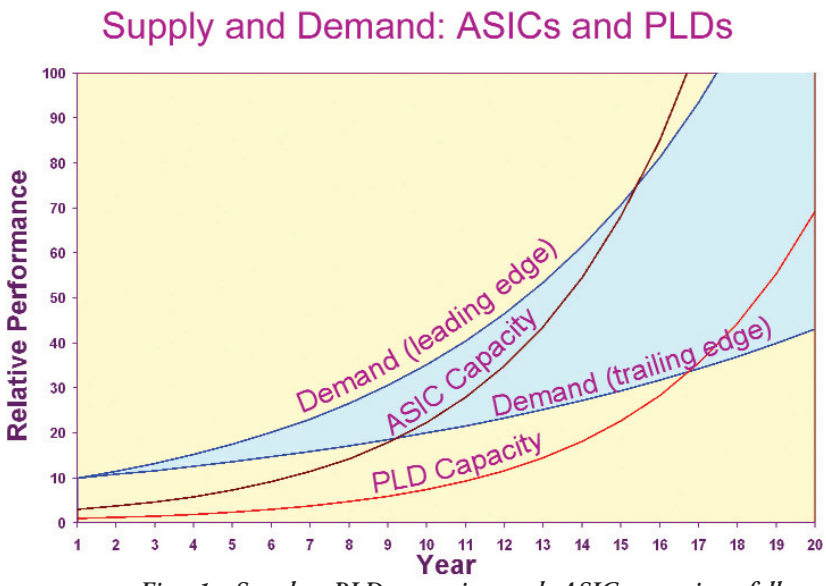

Fig. 1. Supply—PLD capacity and ASIC capacity—follows Moore's law. Demand for capacity grows slower and it spreads across market segments.

Some applications need all the performance they can get; others need all the transistor capacity they can get. Arguments over the merits of ASICs versus PLDs are fought at the leading edge, where ASICs have the advantage. But for the bulk of the market, the requirement is for adequate performance and for adequate transistor capacity. There's still a huge gap in performance and in transistor capacity between leading-edge ASICs and leading-edge PLDs, but PLDs fit the bulk of the market.

ASICs are expensive (high fixed costs), but, for high unit volumes, the per-chip cost can be low. ASICs offer excellent performance. PLDs have low fixed costs, but have high per-chip costs. ASICs are good for systems that require leading-edge performance in high unit volumes. PLDs are good for systems that require low to average performance in low production volumes. There's an opportunity gap between the ASIC's suitability for high-volume systems and the PLD's suitability for low-volume systems. ASSPs and *structured* ASICs are trying to fill this gap.

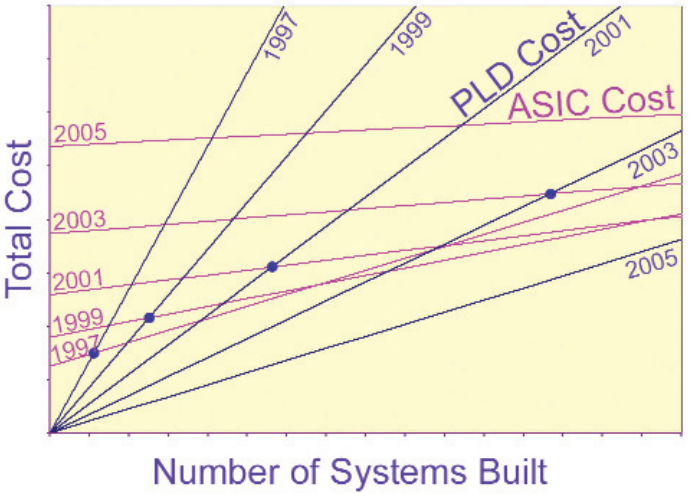

Fig. 2. As ASIC fixed costs rise and as PLD chip costs fall, the equal-cost design point (dots) moves rapidly to the right in favor of PLD-based systems.

"Structured ASICs," from companies such as **Chip Express**, **eASIC**, and **Lightspeed**, provide custom hardware similar to an ASIC, but in a chip that's not quite finished. The customer specifies connections on the last few metal layers. The structured-ASIC provider stocks almost-completed chips in high volume and customizes them with the order. This amortizes most of the design and mask costs (the fixed costs) across a range of applications.

Microprocessors and DSPs

The integrated circuit, combining multiple transistors on a single chip, kicked the semiconductor industry into high gear. It raised the designer's level of abstraction from the transistor to the logic macro. Logic macros—building-block functions made up of multiple transistors—proliferated because they enabled more designers and because they raised designers' productivity.

Ten years of progress culminated in the ultimate logic macro, the microprocessor. Instead of building custom hardware, engineers built a general-purpose system with a microprocessor, memory, and peripheral chips, and programmed it to mimic the behavior of custom logic. These microprocessor-based systems reduced system cost and improved system reliability by using even fewer components than a system built of logic macros. The microprocessor raised designers' level of abstraction from logic macros to programming. Microprocessor-based systems proliferated because they enabled more designers and because they raised designers' productivity.

Enabling more designers and raising the designers' productivity meant giving up something. Compared to custom hardware, microprocessor-based systems sacrifice efficiency and performance. Systems requiring high absolute performance continue to be circuit-based, but for the vast majority of systems, the microprocessor's performance is adequate. The microprocessor market has grown to billions of units a year because a microprocessor-based system is cheaper than a circuit-based system and its performance has been good enough for most applications.

If the microprocessor is buried invisibly in a circuit that mimics custom logic, it is called *embedded*. After ten years of success in embedded applications, the microprocessor was fast enough to be the central processing unit of a computer system. IBM legitimized the Personal Computer in 1981. The PC's introduction split microprocessor design into two camps, a cost-oriented camp for embedded applications and a performance-oriented camp for CPUs (central processing units). The embedded microprocessor camp split further into cost-oriented microprocessors and into performance-oriented microprocessors for signal processing, the DSPs.

CPU microprocessors, selling 150-million units a year at prices 50 to 100 times those of embedded microprocessors, capture more than half of the microprocessor segment's dollars with less than 2 percent of its unit volume.

Problems with Microprocessors and DSPs

After twenty-some years, the PC is good enough. While there will always be users at the leading edge demanding more performance, the PC's performance now satisfies most users. It is supply and demand again; we can recycle fig. 1 by replacing "ASIC Capacity" with "Leading-edge PC" and "PLD Capacity" with "Value PC." Purchases have shifted from leading-edge PCs to "value PCs." Leading-edge PCs offer leading-edge performance at premium prices. Value PCs offer good-enough performance at attractive prices. As the market shifts from leading-edge PCs to value PCs, profit margins will decrease. Engineering resources devoted to increasing the PC's performance will shift to designs for higher-margin untethered systems such as MP3 players, cell phones, and digital cameras. And the microprocessor's design goal is shifting from the cost-performance orientation of the PC's microprocessor to the *cost-performance-per-watt* requirements of untethered systems.

Microprocessors have been supplying performance in a cost-performance segment (the PC), but the microprocessor isn't efficient enough to be the workhorse in untethered systems. Think of the microprocessor as a circuit that constantly reads a user's manual to figure out how to operate itself. The microprocessor's instructions are its run-time operating manual. The microprocessor interprets each instruction and *mimics* the behavior of a custom circuit. It's like trying to perform a figure-skating routine while reading a manual on figure skating. Good luck in the triple jump. Engineers' heroic efforts improve microprocessor performance, but the cost-performance-per-watt environment of untethered systems demands a more direct solution. Microprocessors will move to a supervisory role as the heavy lifting shifts to custom circuits.

PLD futures

The programmable logic companies began in low-end markets by consolidating the system board's logic-macro chips into a single chip that was general-purpose at manufacture and customized in the field. That's the Holy Grail for chips—high-volume production of a generic component with value-added customization by the customer. High volume means lower cost,

TELECOSM TECHNOLOGIES

Note: The Telecosm Technologies list featured in the *Gilder Technology Report* is not a model portfolio. It is a list of technologies that lead in their respective application. Companies appear on this list based on technical leadership, without consideration of current share price or investment timing. The presence of a company on the list is not a recommendation to buy shares at the current price. George Gilder and *Gilder Technology Report* staff may hold positions in some or all of the stocks listed.

Altera (ALTR)

PROGRAMMABLE LOGIC DEVICES

AUGUST 11: 18.85, 52-WEEK RANGE: 8.321 – 20.05, MARKET CAP: 7.218B

Strong company with a twenty-year history of rapid growth, no debt, and plenty of cash.
See this month's issue.

Analog Devices (ADI)

RF ANALOG DEVICES, MEMS, DSPS

AUGUST 11: 36.77, 52-WEEK RANGE: 17.88 – 40.33, MARKET CAP: 13.433B

Reports earnings August 14.

Avanex (AVNX)

ADAPTIVE PHOTONIC PROCESSORS

AUGUST 11: 3.58, 52-WEEK RANGE: 0.63 – 4.95, MARKET CAP: 247.8M

June quarter revenues were just $5.5 million, up from $5.4 in the March quarter. The loss for the quarter was $6.6 million, down from $12.5 million last year. The company halved its R&D expenses from the 2002 June quarter ($3 million versus $6 million), a good sign of financial restructuring but a sure hit to its once paramount place among optical innovators. Avanex now has about $87 million in cash and short-term investments. On July 31, the company completed its acquisition of Alcatel Optronics and certain assets of Corning Optical Components. Revenue and cash arising from these transactions should start appearing in the September financials. The company has said it expects quarterly revenues in the $25 million range and a cash position of more than $200 million.

Broadcom (BRCM)

BROADBAND INTEGRATED CIRCUITS

AUGUST 11: 20.94 , 52-WEEK RANGE: 9.52 – 29.96, MARKET CAP: 5.865B

Broadcom and Intel settled all outstanding patent litigation, which has been going on for years, and agreed to new cross-licensing arrangements. In addition, Broadcom will pay Intel two cash installments totaling $60 million over the next two quarters. Separately, Broadcom announced its GSM baseband processor was selected as the foundation of the highly anticipated Handspring TREO 600, the new version of the popular TREO 300 phone/PDA that should be arriving this fall.

Ciena (CIEN)

METRO WDM PLATFORMS

AUGUST 11: 5.36, 52-WEEK RANGE: 2.41 – 7.74, MARKET CAP: 2.332B

Reports earnings August 21.

Equinix (EQIX)

SECURE INTERNET BUSINESS EXCHANGES

AUGUST 11: 14.549, 52-WEEK RANGE: 2.00 – 20.25, MARKET CAP: 132.2M

June quarter revenue was $28.4 million, an increase of 12% sequentially and 58% year-over-year. Net loss was $21.2 million, but cash used in operations was just $2.7 million, and the company reaffirmed its projection of being cash-flow positive by the September quarter. Equinix has a cash balance of $24.3 million and carries some $170 million in various forms of debt, including a credit facility and senior and convertible notes. The company added a record 87 customers in the quarter, bringing the total to more than 600. They represent more than 90% of the Internet routing table. More than 100 of its customers now participate in GigE Exchange, Equinix's public peering platform, where data traffic grew 30x over the last 12 months and now totals some "tens of gigabits per second." (For reference, 10 Gbps is 3.24 petabytes per month.) Amazon.com also made Equinix its first outsourced data-center provider.

Essex (EYW)

OPTICAL PROCESSORS

AUGUST 11: 5.30, 52-WEEK RANGE: 1.50 – 5.85, MARKET CAP: 47.3M

June quarter revenue was $4.15 million, compared to $729,000 in the year-ago quarter. Net income was $75,000, compared to a net loss of $835,000 last year. Most of the company's sales derive from government contracts: some $2.2 million this year is for work on an advanced radar for the U.S. Missile Defense Agency. Commercial products have yielded sales of $723,000 this year, including $210,000 for five prototypes of the Hyperfine WDM optical telecom module that will soon ship in a more advanced "Alpha" version. The company expects quarterly revenue to mirror the June quarter for the remainder of the year and annual revenue to be $15 million, compared to $4.5 million in 2002.

EZchip (LNOP)

10 GIGABIT NETWORK PROCESSORS

AUGUST 11: 7.29, 52-WEEK RANGE: 3.79 – 8.33, MARKET CAP: 53.2M

EZchip's parent LanOptics (LNOP) reported sales of $423,000 and a loss of close to $3 million. The company spends about $2 million a quarter on R&D and about $1 million on SG&A. It has $18.7 million in cash and marketable securities. Volume shipments of its NP-1c net processor should begin in the fourth quarter, with a concurrent revenue bump expected then or in the first quarter of 2004. Today, its scant revenues

derive from sample chips, testing boards, and software tools.

EZchip was a pioneer in putting memory and logic on the same chip. With an assist from IBM's fab, it was this key insight that catapulted EZchip two generations ahead of the competition. Now the world is catching on, as Taiwan Semiconductor and NEC are building DRAM into logic chips. One report says some 30% of IBM's custom ASIC chips now contain embedded DRAM, though none as memory-rich as the NP-1's 32 Mbits.

As with Telecosm favorites Essex and Corvis, EZchip could get a boost from the U.S. security establishment. The Department of Defense has issued a new requirement that all its global information assets deployed from this October forward be IPv6 compliant. "Internet Protocol version 6" increases each IP address to 128 bits from the current 32 bits used by IPv4, thus expanding the Internet's "address space"—the number of potential connected information appliances—by many orders of magnitude. EZchip is uniquely situated to handle these larger addresses using its current-generation products and has earned design wins in Asia because of this capability. Japan has mandated a national switch to IPv6 by 2005, and China and India, with their combined 2.5 billion people and growing base of IP-addressable PCs and mobile phones, will soon need IPv6, too.

Intel (INTC)

MICROPROCESSORS, SINGLE-CHIP SYSTEMS

AUGUST 11: 23.90, 52-WEEK RANGE: 12.95 – 25.50, MARKET CAP: 155.6B

John Thornton, the former Goldman Sachs president who resigned suddenly last spring to become a professor at Beijing's Tsinghua University, has joined Intel's board of directors. Thornton is on numerous corporate and public-policy boards and gives Intel a robust new communications channel to the top of Chinese politics, business, and academia.

Separately, Intel announced a joint program with Linksys, Cisco's Wi-Fi division, to develop home wireless networks that are easier to set up and configure. Linksys Wi-Fi hubs will sport labels announcing they have been optimized for "Intel Centrino Mobile Technology."

The company also declared a $.02 per share dividend.

National Semiconductor (NSM)

SINGLE-CHIP SYSTEMS, ANALOG EXPERTISE,FOVEON IMAGERS

AUGUST 11: 22.85, 52-WEEK RANGE: 9.95 – 25.24, MARKET CAP: 4.200B

Intel wannabe AMD bought National's Information Appliance division, better known as its line of low-power, low-cost Geode microprocessors. Terms weren't disclosed. The x86-based Geode was always a good idea for areas where the Pentium was clearly an overshoot product, but National could never break through. Now Intel's low-end X-Scale processor is swallowing the appliance markets National hoped to secure. Better to let AMD fight two losing battles with Intel so Halla and Co. can focus on their historical analog expertise and their new line of world-beating imaging solutions, including Foveon. Halla joins us once again at Telecosm in late August.

Qualcomm (QCOM)

CDMA INTEGRATED CIRCUITS, IP, SOFTWARE

AUGUST 11: 35.88, 52-WEEK RANGE: 25.10 – 42.89, MARKET CAP: 28.478B

Encircling its rivals, Qualcomm announced the sampling of the MSM 6250, a single-chip solution incorporating WCDMA, GSM/GPRS, and Qualcomm's popular gpsOne global-positioning capability, radioOne zero-IF technology, and the BREW software platform. Qualcomm has proved itself the leader in both its own proprietary technology, cdma2000, and in advancing the often unworkable technologies of others, such as WCDMA.

Sprint PCS (PCS)

NATIONWIDE CDMA WIRELESS NETWORK

AUGUST 11: 5.33, 52-WEEK RANGE: 1.75 – 6.79, MARKET CAP: 5.455B

June quarter revenue was $3.1 billion, a modest 2.6% increase year-over-year. But churn is down (to 2.4% from 2.9%), and the company says its new customers are of much higher quality than in past years. Sprint signed up 617,000 new mobile-phone users in the quarter, and average revenue per user (ARPU) increased to $62 per month. Sprint's high-speed cdma2000 data offering known as Vision is growing particularly fast. Subscribers to the $10-a-month service, which enables e-mail and picture-mail among other data apps, jumped from 1.3 million in the March quarter to 2.1 million in June.

Synaptics (SYNA)

TOUCH-SENSORS, FOVEON IMAGERS

AUGUST 11: 9.59, 52-WEEK RANGE: 3.13 – 14.90, MARKET CAP: 226.1M

Revenue for the second calendar quarter was $28.2 million, an increase of 8% sequentially and 9% year-over-year. Net income was $2.6 million, or $.10 per share. Sales of non-notebook computer products grew from 2% to 7% of total revenue, a good sign the company is expanding into a wider array of information appliances, and even to automobiles. The stock price, however, plunged 30% on the earnings report, reaching the $10 level not seen since May. The pull-back offers an attractive buying opportunity for a stock up almost 400% from its lows. In addition to its market-leading touchpad solutions, Synaptics is 15% owner of Foveon, Carver Mead's revolutionary one-chip camera company.

Terayon (TERN)

BROADBAND CABLE MODEMS, HEAD-ENDS

AUGUST 11: 4.63, 52-WEEK RANGE: 1.25 – 5.13, MARKET CAP: 341.3M

June quarter revenue was $30.6 million, up 37% from both the previous quarter and the year-ago quarter. The loss was $13.1 million, a number from which the company expects a slight improvement in the September quarter. Terayon has $162 million cash and short-term investments and $65 million in convertible debt.

Xilinx (XLNX)

PROGRAMMABLE LOGIC DEVICES

AUGUST 11: 25.01, 52-WEEK RANGE: 13.50 – 31.00, MARKET CAP: 8.501B

June quarter revenue was $313.3 million, up 3% sequentially. Earnings were $46.2 million, or $.13 per share. Xilinx is demonstrating two of our big themes: the increasing prominence of Asia and the programmable logic paradigm. Japan and Asia-Pac now account for 34% of the company's total sales, up from 25% a year ago; and, consistent with the thrust of this month's GTR, Xilinx's Asian customers are increasingly using its FPGAs to replace ASICs in telecom equipment and even consumer products. The GTR has advocated FPGAs since 1996, and our own Nick Tredennick describes Xilinx and Altera as the next Intels. Sony CTO Dr. Tsugio Makimoto endorses the move of programmable logic into markets previously dominated by ASICs and microprocessors. "Field programmability is a must in this new era of digital consumer goods," says Makimoto. "These products have short time-to-market windows and product life cycles that come quickly and end dramatically. Flexible, agile solutions are critical."

CONTINUED FROM PAGE 3

which makes the chip attractive to more cost-sensitive applications, which increases production volume, and so on.

After a beginning in low-end markets consolidating logic macros, the PLD makers moved into logic-prototyping markets where margins were higher and volumes were lower. As applications blurred the boundary between logic consolidation and mid-range ASICs, the leading PLD companies split their product lines into high-margin, high-capacity chips and lower-margin, lower-capacity chips. For Altera, the high-capacity families are Stratix and Stratix GX and the lower-capacity families are Max and Cyclone. For Xilinx, the high-capacity families are Vertex-II and Vertex-II Pro and the lower-capacity families are CoolRunner-II and Spartan-3. Max and CoolRunner are the lowest-end chips. The Max and CoolRunner families, generally referred to as complex programmable logic devices (CPLDs), feature erasable programmable read-only memory (EPROM) for personalization memory. The other families, generally called field-programmable gate arrays (FPGAs), feature SRAM (static random-access memory) personalization memory.

The low-end PLDs offer capacities in the range of tens of thousands of "gates" to a few million gates. A gate is a logical grouping of transistors—typically six transistors. The cheapest ones are under three dollars. Less than twenty dollars buys a million logic gates. That's a lot of transistor capacity to call the low end!

High-end PLDs offer tens of millions of gates. They also feature memory blocks, complex functions such as multipliers, high-speed communication channels, and even complete microprocessor "cores."

Altera and Xilinx both offer soft-core microprocessors that can be implemented in either low-end or high-end PLD chips. Altera has Nios; Xilinx has MicroBlaze. Nios and MicroBlaze offer PLD designers a ready-made, on-chip state sequencer to manage the application. These soft-core microprocessors offer the advantage of being portable across the family of PLDs and even across generations of products (hard-core microprocessors must be redesigned for each new process generation). In addition, designers can customize these microprocessors to suit a particular application.

With their low-end families, Altera and Xilinx are moving from existing communication and industrial markets to consumer, automotive, and portable markets. With their high-end families, they are moving from prototyping *to markets held by ASICs, ASSPs, microprocessors, and DSPs.*

As they displace ASICs and ASSPs, the PLD companies have an easy task. Engineers building ASICs and ASSPs are logic designers with the same skills and tools required for PLD designs. Many of these projects already build prototypes and initial production units using PLDs. These design teams see PLDs growing into their needs. As costs for PLDs fall, there's less reason to cost-reduce PLD-based systems in the field; it's more profitable to start the design team on the next-generation product.

Invading markets held by microprocessors and DSPs is more difficult. Altera and Xilinx can demonstrate that for signal-processing functions, such as multiply-accumulate (MAC), finite impulse-response filter (FIR), and fast-Fourier transform (FFT), PLD implementations are *tens to hundreds* of times faster than DSP implementations. Not only are they faster, but PLDs run at lower clock rates, meaning that they use less energy. Given the market's need for performance and the PLD's compelling performance advantage, PLD-based implementations should have swept the industry. They haven't.

Problems with PLDs

PLDs use too much energy and they are too slow. As I said, PLDs are wires, logic blocks, and personalization memory. Wires and personalization memory dominate the chip. Because of the personalization transistors and because of the extra wiring, connections are slower than directly wired connections, so the chip is slower than its ASIC alternative. General-purpose PLDs provide many options for connecting wires and logic blocks and each SRAM personalization bit uses six transistors. As long as there is this huge overhead in wires and in personalization memory, there will be a wide gulf between the performance and capacity of leading-edge ASICs and the performance and capacity of leading-edge PLDs. This gulf is easily seen between the ASIC and PLD capacity lines in fig. 1.

Help is on the way to reduce wiring overhead: 3D wafer stacking. The industry has been so focused on reducing transistor size (because that worked and manufacturers knew how to do it) that it has been blind to options that may be more cost-effective than shrinking transistors. One option stacks wafers with vertical wires connecting between layers. It is possible to stack a dozen or so wafers with *hundreds of thousands* of vertical connections between chips. *This makes it possible to mix digital logic, memory, and analog circuits—each designed in its own optimum process.*

IBM, **Tezzaron Semiconductor** (formerly Tachyon Semiconductor), **Ziptronics**, and others are experimenting with 3D chips cut from 3D wafers. Stacked chips have dramatically reduced delays due to wiring, because wires run a few *microns* (thousandths of a millimeter) vertically rather than running *millimeters* across a 2D chip. Manufacturers can even stack redundant circuits or memory to improve yield. Thousands of vertical copper wires connecting stacked wafers prevent heat accumulation in internal layers. Stacked wafers are thinned. This aids heat removal and it means that the final stacked chips can use *standard packages* because they are not noticeably thicker than 2D chips.

The PLD's conceptual model stays the same. It's general-purpose interconnect is much more efficient in three dimensions than it is on a 2D surface. Wires would no longer dominate chip performance.

Help is also on the way to reduce the overhead of the

PLD's personalization memory. The market shift to untethered systems demands the ultimate in memories: non-volatile like flash memory, dense as DRAM, and fast as SRAM. Flash memory, DRAM (dynamic random access memory), and SRAM—all developed for the PC—have crippling disadvantages in untethered systems. Flash is very slow and it wears out. DRAM is slow. SRAM uses too much energy. And DRAM and SRAM lose their memory when the power is off.

Untethered systems have empty sockets to fill, waiting for the ultimate memory chip. That's the investment incentive that will finally get non-volatile memory to market. Long-time candidates, stalled in market growth by the incumbents' lock on the PC, may soon gain a foothold. These candidates are FRAM (ferroelectric RAM), MRAM (magnetoresistive RAM), and OUM (ovonic unified memory). FRAM stores bits in electrically polarized crystals, MRAM stores bits in magnetic domains (similar to today's hard-disk recording), and OUM stores bits in amorphous or crystalline states (similar to a CD or DVD recording). But newcomers, such as programmable metallization cell memory (PMCm), have a good chance too.

A non-volatile memory cell will greatly improve the density, speed, and security of SRAM-based PLDs. Security improves because configuration bits are held on the chip and do not travel across the chip's interface each time the chip is powered up. PLDs with non-volatile memory will compete better against ASICs, ASSPs, microprocessors, and DSPs and they will be efficient enough for untethered applications.

Reconfigurable

Let's start with *configurable*. **ARC International** (ATVL.PK) and **Tensilica** offer configurable microprocessor cores. This is a step in the right direction. With most microprocessors, designers take what they can get from the manufacturer. There may be no microprocessor that has the perfect instruction set for a flatbed scanner, for example. But the engineers at the scanner company can specify special hardware and special instructions for an ARC or Tensilica microprocessor that will greatly accelerate the company's proprietary scanning algorithms. Development tools spit out a custom microprocessor together with appropriately modified operating systems, compilers, assemblers, and test code.

The next step in the development of configurable processors will be for ARC or Tensilica to accept an executable specification and to spit out a black box and an application program. That's a path to success because configurable processors offer more performance than fixed-instruction-set microprocessors and they still enable programmers. ARC and Tensilica offer a smooth transition from today's microprocessor-based designs to the future's circuit-based designs, a transition that takes the programmers along.

But configurable microprocessors aren't *re*-configurable. To meet my definition of reconfigurable (there is no standard definition), the system's structure must change while the machine is running. That's something a standard microprocessor, or even a configurable microprocessor, does not do. General-purpose PLDs could do this, but they aren't built for that because the large PLD manufacturers have been growing with and following their customers in logic consolidation and in prototyping. For those applications, there's no requirement for partial reconfiguration or for run-time reconfiguration. Consequently, today's general-purpose PLDs don't support rapid reconfiguration or partial reconfiguration well.

Reconfigurable systems are a hot topic. Reconfigurable is popping up everywhere—as a marketing label for conventional systems. Real reconfigurable systems are still rare. The advantages of reconfigurable systems are in resource availability and in adapting to change. Startup **QuickSilver Technology** builds reconfigurable systems.

The difference between PLDs and reconfigurable chips begins with design approach. PLDs are bottom-up; QuickSilver's approach has been top-down. The logic blocks on PLDs developed from logic macros. Instead of beginning with logic macros, QuickSilver's engineers began with the algorithms. They then broke the algorithms in the application set into their fundamental constituents. QuickSilver's chips are arrays of these algorithmic "nodes."

A reconfigurable system builds custom structure for

Microprocessor Performance

The microprocessor's performance has been improving at 60 percent per year. Much of that performance comes from increasing the clock rate. The Intel 8088 microprocessor that drove the original personal computer ran at 4.77 MHz. Leading-edge PC microprocessors in 2003 run at 3,000 MHz. That's a clock-rate increase of 34 percent per year. The rest of the performance comes from branch prediction, wider data paths, parallel function units, pipelining, caches, and a host of other hardware enhancements.

The PC market has dominated performance-oriented microprocessor applications for the last twenty years. The PC is a performance-oriented consumer item, so the PC microprocessor's design objective has been cost performance. But increasing the microprocessor's performance by increasing its clock rate increases power consumption. Doubling the clock rate doubles the power consumption.

The clock rate of today's leading-edge microprocessors is 600 times the clock rate of the PC's original microprocessor. If its energy use rose by the same amount, the microprocessor wouldn't be usable. Therefore, manufacturers lower the microprocessor's supply voltage. Cutting the voltage in half lets the microprocessor run at four times the clock rate using the same amount of energy.

Microprocessor makers have traded almost all the available voltage to get to today's high clock rates. Microprocessors running much below one volt approach the operating limit of their transistors.

processing at a particular time. It might begin by configuring a custom structure for encryption, then, using the same resources, it might configure a custom structure for a fast-Fourier transform. Think of the chip as a pool of physical resources that can be structured as needed at run time.

The challenge

As an engineering discipline, the problem with the microprocessor and the DSP is that they have stalled progress in design methods for thirty years.

Digital systems have two defining characteristics: structure and procedure. The structure is what and the procedure is how. In microprocessor systems, the microprocessor provides the structure and the engineer provides the procedure in the form of programs. In automotive systems, the automobile is the structure and driving is the procedure. There are many more drivers than automobile designers, but someone has to design the engines, brake systems, seats, headlights, drive trains, and so on for the rest of us. For most of us, the level of abstraction in automotive systems is driving. The same is true for the community of design engineers, the level of abstraction in system design is programming. A few engineers design the microprocessors that the rest of the design community uses.

That isn't a problem in automotive systems. But untethered systems will dominate the future and the microprocessor's days are numbered as the workhorse in such systems.

In the days of logic macros, the designer was responsible for both the structure and the procedure. With the microprocessor, few engineers built microprocessors (structure). Most designers chose the structure for implementation (microprocessor) and built systems by providing procedure in the form of programs. In thirty years of building microprocessor-based systems, *problem-solving has become programming*. That's what the design community knows, that's what the universities teach, that's what the installed base of development systems supports, and that's the model microprocessor manufacturers work to sustain.

Designing PLD-based systems is much more like designing with logic macros than it is like programming a microprocessor. PLD makers can invade the market for ASICs and ASSPs easily because they are invading a market that requires the same design skills. In attacking the microprocessor and DSP markets, the PLD makers will be unsuccessful unless their chips and design procedures are accessible to programmers who are not logic designers. That's both an enormous challenge and an opportunity. Companies like **AccelChip**, **Celoxica**, and **MathWorks** have products that raise the level of abstraction in PLD design to programming.

We are watching the standard recipe for disruption play out. The PLD, with inherent advantages, is coming in under the microprocessor and ASIC makers' radar. It's the replay of a theme we've seen before (workstations and PCs, leading-edge PCs and value PCs, leading-edge transistors and value transistors) with analogous arguments and forecasts. ASICs are the workstations of chips. Initially, 100 percent of demand is at the leading edge. Over time, demand spreads into market segments. Meanwhile, supply increases with Moore's law. Low-end suppliers enter the market. The market leaders abandon the low end, working to keep their leading-edge customers. Eventually, building for the leading edge becomes inertia as leading-edge producers leave their customers behind. The low-end suppliers build for volume and ride Moore's law to performance, capturing the bulk of the market.

The PC came in under the workstation makers' radar. The general-purpose PC had few features and dismal performance. All it had going for it was mass production for consumer markets. The performance of PCs and of workstations grew faster than demand for performance rose. Workstations built for leading-edge demand and outgrew most customers' needs. The PC built for volume and grew into the spreading demand for performance. The PC's volume-based strategy trounced the workstations' performance-based strategy because it rode volume to low cost and it rode Moore's law to performance.

PLDs are coming in under the ASIC makers' radar. General-purpose PLDs have fewer transistors and have lower performance than leading-edge ASICs. The capacity and performance of PLDs and of ASICs is rising with Moore's law. Demand is rising slower. ASIC makers are following their leading-edge customers and will exceed the needs of most of the market as PLDs move in.

The objective has always been finding the right generic, high-volume component. The transistor was the first generic, high-volume semiconductor; the microprocessor was the second. Next it will be the PLD. This is evolution.

—Nick Tredennick and Brion Shimamoto
August 11, 2003

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@GILDER.COM

EDITOR IN CHIEF	EDITOR	EXECUTIVE EDITOR	TECHNOLOGY ANALYST	MANAGING EDITOR	ART DIRECTOR	SUBSCRIPTION DIRECTOR	CUSTOMER SERVICE
George Gilder	Nick Tredennick Brion Shimamoto	Bret Swanson	Charlie Burger	Mary Collins Gorski	Peter Johnstone	Rosaline Fernandes	Sandy Fleischmann

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