GILDER TECHNOLOGY REPORT

July 2006 / Vol. XI No. 7

Semiconductor Outlook

I'm sticking my neck out to say where semiconductor technology is headed in the next ten years. I start with trends that are obvious to everyone: device mobility, ubiquitous access, and emerging economies. From there, I explore causes, consequences, and derivative trends. Last, I say what the trends mean to specific market segments and companies.

Device mobility. My brother, John, builds houses. In the old days, he had three choices at a new job site. He could wait for the local utility's line power installation, he could bring a generator, or he could work with unpowered tools. Today, he uses battery-powered hand tools that are powerful and efficient enough to last all day. The convenience of portable devices that is compelling for construction work applies across a wide range of occupations and activities, extending mobility to toothbrushes, hair dryers, phones, laptops, PDAs (personal digital assistants), and remote controls.

(CONTINUED ON PAGE 3)

For patient investors focused on the next five to ten years, Altera and Xilinx beckon as anchor holdings for a high tech portfolio.

Inside:

The FPGA Tourney: Altera and Xilinx

Now is the time to invest in programmable logic devices (PLDs), and there's no better way to do that than through **Altera** (ALTR) and **Xilinx** (XLNX). Here are six reasons why.

Shrinks accelerate growth

In contrast to ordinary chips with designs permanently fixed in manufacture, PLDs are customized in the field by the users, who interconnect or program a mesh of logic circuits by blowing or retaining fuses between prefabricated "gates." Users thus bypass the complex and costly upfront procedures used to make and test application-specific integrated circuits (ASICs).

The downside of programmability, however, has always been increased chip size and exponentially higher costs. Larger chips are more likely to be defective and fewer of them fit on the wafer. A 10% increase in size, for example, might double a chip's cost. To lower these costs, customers often design prototypes using PLDs, particularly field-programmable gate arrays (FPGAs), that can be programmed in their sockets, and then transfer the design to an ASIC for volume production.

As shown by the downward swoon of ASIC sales, this pattern is now breaking down. As Moore's law doubles the number of transistors on a chip every 18 months, the up-front costs rise for design, verification, and masks for ASICs. While becoming larger and targeted for more specialized niches, ASICs need ever-larger volumes to make them economical.

Meanwhile, the same Moore's law trend improves the cost, speed, and capability of FPGAs that can be programmed after production for thousands of different customers. This means ASICs will be marginalized to the highest volume products or to bleeding-edge applications where performance is more important than cost. FPGAs ultimately take over. As we said last month, the industry could end up making just one Altera or Xilinx chip with nearly infinite volume and make it available nearly free. You can pay for the software that programs it.

Flexibility drives designers fast forward

Further reducing ASIC volumes is the acceleration of technology, making products obsolete faster. To speed

development and reduce design costs, FPGAs can be reprogrammed to upgrade products or to correct for design errors.

Shortening FPGA design time are design blocks known as intellectual property cores (IP cores), which embed standard functions into programmable logic chips using hardware macros and extensive libraries of soft macros. IP cores allow logic programmers to focus on designing their unique systems rather than spending time designing off-the-shelf functions. Among the most popular on-chip macros are microprocessor cores, preferred because there are ten times more programmers than logic designers and because programming is a higher-level abstraction than logic design, making microprocessor engineers more efficient. With processors on chip, FPGAs can invade the huge markets for microprocessors.

Altera and Xilinx own the market

The reigning giants of programmable logic for over two decades, Altera and Xilinx together own about 87.6% of the growing market, with thousands of customers in communications (networking, wireline, wireless), computers and storage, consumer products, and industrials (automotive, instrumentation, military, security, energy, management). Their market share has steadily increased from 81.9% since 2001, while laggers **Lattice Semiconductor** (LSCC) and **Actel** (ACTL) have steadily lost share, from 18.1% to 12.4%.

With combined operating income of over $700m, Altera and Xilinx earn virtually all the profits in PLDs while Lattice suffers chronic losses and Actel waffles about breakeven by several million dollars. Focusing on the niche markets of military, avionics, and space-grade FPGAs, Actel's highly secure and radiation hardened chips can only be programmed once, severely restricting their commercial applications.

For these seeking to profit today from the coming boom in FPGAs, Altera and Xilinx are it. Both companies are financial and operational powerhouses. Since the most recent economic downturn in 2002, each has improved its market share, with Xilinx nudging up to 52% from 51.1% based on sales of $1,645m during calendar year 2005, while Altera grew its share to 35.6% from 32.4% on sales of $1,24m. Xilinx ended fiscal year 2006 in March with $403m of free cash flow, $1,719m of net cash, and no long-term debt. Similarly, Altera ended fiscal 2005 in December with free cash flow of $389m, net cash of $805m, and no debt.

PLDs are a sticky business

An FPGA family typically reaches peak sales 4 to 5 years after introduction, and designs introduced during 2002–05 are still ascendant. For Altera, that includes the high-end Stratix family for high-capacity applications and the low-cost Cyclone line for high-volume applications such as digital set-top boxes, DVD player/recorder systems, automotive telematics, and flat-panel televisions. Finally, the Max line consolidates miscellaneous logic, called *glue logic*, that ties microprocessors, memory, and peripherals together into a working system. Personalized only once, Max chips are low cost and high volume. The Xilinx families that correspond to Max, Cyclone, and Stratix are CoolRunner, Spartan, and Virtex, respectively.

Stratix and Cyclone have driven Altera's revenue growth over the past two years. During the March quarter, Altera's new product sales rose 173% over the year-ago quarter while older "mainstream" products grew 35%. Over the same period, sales of Xilinx's new products grew 103% while mainstream products declined a slight 3.8%. Introduced last fall, the Virtex-4 family now generates more in 90-nanometer sales than any other PLD line; Spartan-3 continues to lead in number of 90-nm units shipped. Last month, Xilinx begun shipping 65-nm products, putting it ahead of Altera at that process geometry by about half a year.

Seesaw rivalry incites innovation

As a customer becomes familiar with a particular vendor's programming software, provided nearly free, he tends to use it to create other designs. This gives that vendor an edge when competing for the customer's future business. A major Altera software upgrade released in early 2001 eliminated a serious deficiency in this field, and Altera is now well positioned with its refreshed product line; the company claims to be the leading supplier of FPGAs introduced since 2002.

In addition, Altera uniquely offers a line of structured ASICs called HardCopy. In contrast to traditional ASICs, in which customers design every mask layer, Altera's structured ASICs share a common set of base layers developed from equivalent FPGAs. Designing only the last few mask layers, customers gain near ASIC performance with lower development costs, reduced die size, and shorter lead-times.

While both companies offer soft IP cores, Xilinx also offers chips with from one to four hard PowerPC cores. Compared to soft cores, Xilinx's hard microprocessor cores perform better in less silicon area but create more complexity in the product line. By contrast, Altera recently announced a new microprocessor productivity tool for its Nios II soft core that enables, for the first time, software engineers to design hardware. Nios is already the industry's most widely used soft microprocessor.

It's an order of magnitude opportunity

Displacing fixed-chip logic will be a huge deal for programmable logic vendors over the coming decade. For example, during the most recent quarter, Altera reported design-win bookings that were 25% higher than its previous quarterly record, fueled by ASIC and ASSPs (application-specific standard product) replacement designs. Based on public data and information derived from Gartner Dataquest, the PLD market was about $3.3b in 2005, whereas digital logic sales, primarily ASICs and ASSPs, came to some $34b.

Investors seem to grasp neither the enormity of the Moore's law opportunity nor the titanic technology and financial leads enjoyed by both companies. With the recent dip to $22.31, shares of Xilinx are selling for just 21 times the backward-looking free-cash-flow for fiscal 2006 that ended in March. Yet more remarkable, at $17, Altera is on a fire sale at 16 times free-cash-flow for fiscal 2005 ending December. There remains huge room for growth. If revenues for both companies double and ASIC sales stagnate, Altera and Xilinx would still barely ascend to 16% of the ASIC market (while the microprocessor market, an order of magnitude larger, beckons for the next decade). For Xilinx, a conservative projection would mean $2.26 per share of free cash flow at today's margin of 23.3% for a share price of $56.50 at a 25 times multiple over free cash flow. Better yet, Altera's price would triple to $51.75 based on its more profitable free-cash-flow margin of 34.6% and a 25 times multiple.

For patient investors focused on the next five to ten years and unfazed by silly pettifoggery over options, Altera and Xilinx beckon as anchor holdings for a high tech portfolio.

— Charlie Burger, June 28, 2006

(CONTINUED FROM PAGE 1)

Ubiquitous access to voice, video, and data is almost an extension of device mobility. Not so long ago, I visited a library for information, made calls when I got to a convenient phone, and played music or watched movies at home. Today, I get frustrated when I can't answer a question immediately using the Internet, when I can't make or take a call anywhere, and, probably soon, when I can't see or hear whatever I want at any time or place—whether it's about Civil War battlefields or product and pricing information about items in a nearby store. And I expect the information to be free. Such access will be so pervasive that everyone will take it for granted.

Emerging economies. Humans, as Julian Simon demonstrated, are the cylinders of the world's wealth-creation engine. The world economic engine has been running on less than a third of its cylinders. With China, India, Eastern Europe, and others joining the world economy, that engine of wealth creation will soon engage two thirds of the world's population. The world's producers will scramble to meet the demands of emerging economies for household appliances and for other goods that the developed world takes for granted.

Consequences of device mobility

Proliferation of mobile devices will drive advances in power sources, in memory cells, in antennas, in microelectromechanical systems (MEMS), in sensors and actuators, and in analog and digital circuit implementation. The four major areas of power-source development are small engines such as microturbines, fuel cells, batteries, and supercapacitors. I haven't studied this market, but it's a challenge and there's plenty of room for improvement.

Memory cells. The current memory types—SRAM (static random access memory), DRAM (dynamic random access memory), and flash memory—grew up with the PC, which exploited their strengths as it masked their flaws. None of these memory types is suitable for mobile applications. SRAM and DRAM lose their contents when power is off. SRAM is a power hog. DRAM is slow for reading and for writing. Flash memory is slow for reading, very very slow for writing, and it wears out. Mobile devices need a memory cell that has the speed of SRAM, the density of DRAM, and the non-volatility of flash memory. The growing market for mobile devices provides the financial incentive to develop this "Holy Grail" memory cell. Candidates include ferromagnetic memory (FRAM), magnetoresistive memory (MRAM), ovonic unified memory (OUM), and a host of "nano"-based memory types, such as those from **Nantero** and **Zettacore**.

Four years ago, I wrote a *Dynamic Silicon* issue (May 2002) assessing the three non-volatile memory candidates. It was too early to guess the winner, but I favored OUM. At the time, the candidates were announcing prototypes or products and commercial success a few months to, at most, two years away. Unfortunately, concepts and prototypes are easier to demonstrate than to manufacture in volume, and after four years we still don't have a successful product (though OUM from **Energy Conversion Devices** (ENER) still seems most promising).

In general, new stuff may use unique or at least non-standard processes, and it has very low volume, so costs are high. The industry will pay a premium for new memory cells unavailable elsewhere, but to capture even a tiny fraction of the memory market, the maker will have to deliver millions of components. So far, no one has been able to do it. Meanwhile, producers of traditional memory cells are improving their own technology, raising the bar for the innovators.

But the incentive remains, so expect the problem to be solved in the next few years, with perhaps several niche solutions that will, collectively, displace the incumbents and capture the mobile market.

The new memory cell will ripple through programmable logic companies such as **Altera** (ALTR) and **Xilinx** (XLNX). SRAM-based FPGAs (field-programmable gate arrays) just aren't efficient. Imagine replacing a six-transistor memory cell with a single transistor and you get an idea of the benefit to FPGAs in lower cost or in higher capacity. Chips based on a non-volatile memory cell would be faster, cheaper, denser, and non-volatile cells don't need to be configured each time the power comes on.

But the big producers, Altera and Xilinx, may well not dominate the technology. If there's the slightest problem with the new memory process, the rival would tear the experimenter apart. So the likely initial beneficiary will be a startup such as **Ascenium**, **Element CXI**, **KLP Logic**, or **Tabula**.

Circuit implementation combines electronic circuitry with design goals. The move from wall socket power to mobile power causes a change in the design goal from cost-performance to cost-performance per watt, and this change ripples through the entire development process.

While most of an FPGA is digital, analog circuits drive the subsystems closest to the radio's antennas and to the device's sensors and actuators.

Digital circuit implementation (logic implementation) falls into two major categories: microprocessor-based design and custom hardware. Custom hardware is sometimes called ASIC design or SoC (system-on-chip) design.

With time, rising mask and development costs push ASICs toward larger markets, while increasing microprocessor performance and decreasing costs at the low end grow microprocessors in all directions, increasing their overlap with ASICs. Both trends decrease oppor-

TELECOSM TECHNOLOGIES

- Advanced Micro Devices (AMD)
- Altera (ALTR)
- Analog Devices (ADI)
- Broadcom (BRCM)
- Broadwing (BWNG)
- Cepheid (CPHD)
- Corning (GLW)
- Energy Conversion Devices (ENER)
- Equinix (EQIX)
- Essex (KEYW)
- EZchip (LNOP)
- Finisar (FNSR)
- Flextronics (FLEX)
- Ikanos (IKAN)
- Intel (INTC)
- Microvision (MVIS)
- National Semiconductor (NSM)
- NetLogic (NETL)
- PMC-Sierra (PMCS)
- Power-One (PWER)
- Qualcomm (QCOM)
- Semiconductor Manufacturing International (SMI)
- Sigma Designs (SIGM)
- Semitool (SMTL)
- Sprint Nextel (S)
- Synaptics (SYNA)
- Taiwan Semiconductor (TSM)
- Texas Instruments (TXN)
- Xilinx (XLNX)
- Zoran (ZRAN)

Logon with your subscriber ID at www.Gildertech.com for company updates and analysis.

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.

Ikanos (IKAN)

PARADIGM PLAY: VDSL PIONEER

JUNE 28: 14.14; 52-WEEK RANGE: 9.36 – 24.97; MARKET CAP: 383.52M

Like **Sigma Designs** (SIGM), Ikanos is suffering from margin erosion along with inventory builds and growing competition. But, competition is actually good news for this fabless VDSL pioneer (very high- speed DSL). To date, virtually all the VDSL ports shipped worldwide have been supplied by Ikanos, giving the company a huge early advantage in rolling out its next-generation VDSL2 products. But systems vendors and carriers frequently shy away from single-source technologies, and the emergence of competitors will give them the confidence that the industry is ready to support serious VDSL2 deployments, which CEO Rajesh Vashist believes will begin later this year. Bracing for these deployments, systems houses may have been building inventories. Thus, Ikanos believes its sales growth may moderate a bit over the next few quarters while expenses continue to increase due to ramp-up costs for high-margin chipsets, which are expected to begin shipping later this year.

All these factors have been pressuring earnings from operations (excluding options expensing and noncash acquisition costs), which fell from $0.20 per share in December to $0.14 in March to a likely $0.06 in June based on our interpretation of company guidance. The earnings graph looks ugly, and by the time of our 3 May analysis it had incited investors to gouge a third out of the stock's mid-February high of $24.97. Down 43% at the recent price of $14.14, investors are either anticipating a coming collapse of broadband and the Internet or believe that Ikanos is secretly planning to self-destruct.

If Ikanos achieves its reasonable long-term model of 15% operating margin and 20% tax rate for 2008, sales would need to degenerate to $118m if the company is to earn its current price at a growth PE of 30. (EPS would rise more than 30% per year in this case.) That's absurd. Revenue of $118m is notably below annualized sales of $160m anticipated this quarter and a mere 23% of the expected half billion dollar market for VDSL silicon in 2008. Plus, there's plenty of upside potential here. For instance, to earn a double from today's price, at a PE of 30 Ikanos need only capture 46% of the 2008 market.

Not content to rest on its near monopoly of this rapidly growing market, Ikanos continues to innovate and integrate, and has recently added to its product portfolio wire-speed NPUs, VoIP (voice over IP) and security engines, and wireless-LAN (local area networks) capabilities through its acquisition of **Analog Device**'s (ADI) gateway products. The ADI suite will enable Ikanos to offer, in addition to chipsets used for layer-1 functions in central office DSLAMs (DSL access multiplexers), silicon solutions for full-featured gateways targeting both homes and businesses.

Targeting your telecosm portfolio at this bargain should be IKAN.

National Semiconductor (NSM)

PARADIGM PLAY: ANALOG LEADER AND IMAGER PIONEER

JUNE 28: 22.49; 52-WEEK RANGE: 21.24 – 30.93; MARKET CAP: 7.56B

You need analog to link to the world, which is why the world needs National Semiconductor (NSM) with its high-speed interface and analog paladin for "gigabits and milliwatts."

By focusing on precision and high-speed products with the lowest possible power consumption, National continues to increase its share of the standard linear market. During the first 11 months of fiscal year 2006 ending in May, the company's standard linear sales grew 20% while the global standard analog market rose just 12%. National grew faster than the industry in all four product categories—interface (up 50%), data converters (up 40%), and amplifiers and power management (each up 25%).

In the latest quarter, standard linear products garnered 74% of company-wide sales compared to 72% a year ago. Though relatively new areas for National and still only 15% of revenues, data converters and interface products improved hugely in the quarter, with data converters up 30% sequentially and interface up 10%. Both lines sport high margins and helped drive gross margin to a record 61.4% from 54.7% a year ago. Even more impressive, operating margin jumped 11 points over last year to 34.5%.

And … revenue *rose* 12.8% compared to last year's quarter *despite* the pruning of nonaligning units such as the Advanced PC division with its digital and mixed-signal IP, the cordless phone business in Europe, and the Singapore test and assembly plant, which specialized in high pin-count digital packages. (A newly opened test and assembly facility in Suzhou, China, supports National's analog business along with a facility in Asia.) As low margin foundry sales continue to fade out over the next few quarters, gross margin should bulge past 65%, though revenue may come under slight near-term pressure.

Here's more good news for investors—National's bull run has only just begun. Back when PCs were driving

Online Bonus Material: For additional analysis on **Analog Devices (ADI)**, **Essex (KEYW)**, and **Semitool (SMTL)** logon with your GTR subscriber ID at *www.Gildertech.com*.
Visit *www.TelecosmConference.com* today to register for Gilder/Forbes Telecosm 2006, October 4 - 6 in Lake Tahoe, and receive half-off the Forbes.com registration price.

demand, semiconductors made up 10% to 20% of device content. Today, semis comprise 40% of the content of the products propelling demand. Increasingly that content is analog, particularly standard linear for enhanced displays, improved audio, better wireless, and longer battery life. Needed in every flat panel screen and in every high-resolution graphics systems will be National's low voltage, differential signaling devices that function in low noise, low power analog at over 2 gigahertz speeds. The high-definition heyday for this technology looms ahead even while China emerges with its 400 million new consumers with cell phones and other gadgets, in addition to other emerging nations consuming at an accelerated pace.

The standard linear market is now growing some 15% annually and CEO Brian Halla expects National to ascend significantly faster. His long-term goal has standard analog sales approaching 100% of the company's total revenue. Shorter term, standard linear is on track to exceed 80% of revenue by next year.

National's finances have already responded to Halla's prescription: Free cash flow (cash flow from operations minus capex) for fiscal year 2006 increased 47% over fiscal 2005, to $633m. The $201m increase was a full 82% of the $245m fiscal year revenue increase. Free-cash-flow as a percent of revenue increased from 22.5% to 29.3%. Thus, at the recent share price of $22.49, National was on a sale at a meager 12.5x this backward-looking cash-flow margin. For a paradigm ascender with swelling markets and margins, with cash and receivables of $1,250m (3x current liabilities), with virtually no long-term debt, and with capex expected to remain below 10% of rising revenue, that's a buy signal with a capital B.

Sigma Designs (SIGM)

PARADIGM PLAY: LIFE AFTER TV MEDIA PROCESSOR LEADER

JUNE 28: 9.16; 52-WEEK RANGE: 7.40 – 17.05; MARKET CAP: 208.79M

Sigma Designs has had another 26% trimmed off its market cap since we warned you, just four weeks ago, to expect more short-term softening of the stock before a serious ascent later this year or early next. A dollar of the drop came during the final two hours of trading on 26 June, when a mysterious buzz cut shaved 9% off the stock. Who took Sigma to the barber? Speculation ranges from technical traders to conspiratorial short sellers to inept company accountants to greedy backdaters. But all you need to know is that the bargain just got better.

Probably the major reason for the Street's longer-term trim was the dramatic drop in gross margin, which plunged from 69% last October to 60% during the January quarter on its way to 52% last quarter as Sigma's initial premium prices succumbed to aggressive high-volume pricing. Despite accelerating top-line growth, therefore, operating income has continued to hang near break-even over the past year. Worse for traders fixated on company miniscules, gross margin is expect to fall yet further to 49% this quarter before edging back to the long-term target in the low 50s during the second half of the year as Sigma wrings out chipset launch costs while benefiting from volume shipments and yield improvements.

Overlooking the good news—that these expenses and challenges are typical of chipset ramps, that volume orders mean market takeoff and consummated marriages with sticky customers—skeptical traders instead appear to have blinked at the forward PE of 70 that we estimated in our previous analysis, which was based on the 1 June price of $12.54 and management's conservative revenue forecast for the fiscal year ending next January. Also lost on the cynics is the life-after-television paradigm and the lifeline of opportunities it means for Sigma. By plunging into the market early and putting all the standards on one chip, Sigma has achieved a real edge against **Broadcom** (BRCM), **Texas Instruments** (TXN), **Intel** (INTC), and others who assumed they could wait until the standards were settled. The company's media processors are set to litter the landscape of the coming IP video world, including IPTV set-top boxes, high-definition DVD players, portable media players, and flat panel TVs.

Now flirting with $9.25 and *still* sporting a seemingly stratospheric PE of 51, the stock presents long-term investors with near cast-iron protection on the downside—Sigma need only nudge revenue up 15% to $85m next fiscal year in order to increase earnings per share (EPS) by 51% and earn its current multiple. And that's off a revenue estimate for this year that's likely much too conservative; based on current design wins and orders, Sigma could well achieve those sales *this* year. To stumble at its current valuation, the world would have to turn off all its flat screens and media players.

On the upside (also off our conservative outlook for this year), if Sigma increases sales another 121% next year—doable based on the stickiness of design wins and the company's significant technology lead—$48.6m settles to the bottom line for an operating EPS of $1.62. With a growth PE of 30, the stock quintuples.

— *Charlie Burger*

tunities for ASIC design companies such as **Fujitsu**, **IBM** (IBM), **LSI Logic** (LSI), **NEC** (NIPNY), and **Toshiba**.

Estimates of market size for ASICs and microprocessors differ substantially. The Semiconductor Industry Association forecasts a semiconductor market for 2006 of $245 billion with growth to $309 billion by 2008. Of that, the microprocessor market is $40 billion for 2006 and is projected to grow to $46 billion by 2008. The microcontroller and DSP (digital signal processor) markets are $13 and $9 billion, respectively, and are expected to grow to $16 and $13 billion by 2008. Traditionally based on the predictable growth rates of demand in developed countries, these estimates may well have to be revised upward to address rates of growth in emerging economies. In all estimates, the ASIC market is smaller and is growing slower than the microprocessor market.

With devices that plugged into wall sockets, power-efficiency didn't matter as much as cost-performance. In general, minimizing product cost meant minimizing development and bill-of-materials costs. Making engineers more productive reduces development cost; using generic components, such as microprocessors, standard peripheral chips, and memories, reduces bill-of-materials cost. Reducing engineering cost meant raising the level of abstraction from logic design to assembly programming and then to high-level-language programming. Reducing bill-of-materials cost meant building with microprocessors and memories. But using high-level languages and microprocessors reduces power-efficiency because the microprocessor and memory combination simulates functions that would otherwise be streamlined in custom hardware.

After 35 years, microprocessor-based design is the standard, but it isn't power-efficient for mobile devices. The alternative to microprocessor design, custom hardware, is too expensive to develop and it doesn't accommodate rapid product cycles or fast-changing protocols that are characteristics of consumer mobile devices.

Moreover, software programmers can build microprocessor designs, but custom hardware requires the skills of logic designers. Programmers outnumber logic designers ten to one, so the future lies in building systems with programming skills. That's a big transition for companies like Altera and Xilinx with a culture of circuit design and logic design.

Founded by circuit designers, these companies evolved one step up into logic design as programmable logic complexity increased. Today they offer some programmer-oriented software development tools, but the primary orientation of their software systems still supports logic designers. The many customers of Altera and Xilinx are encouraging them to continue business as usual. It will work for some years, but programmer-oriented applications, often from new companies, will eventually outgrow logic-oriented applications, which will then begin to decline just as ASIC design is now or soon will be declining.

Emerging is a combination of custom logic and microprocessor-based design. But because the new regime will support programmers as the application engineers, it will grow out of the microprocessor application space rather than out of the ASIC space and the design interface will be high-level languages familiar to the systems engineering community. The microprocessor will be the core because the essence of the microprocessor is the state sequencer that embodies the intelligence in every system. But, because microprocessor-based design isn't power efficient, the engineer's high-level-language description will be turned partially into custom hardware. The custom hardware will be "soft" in that it will be paged into a generic component such as an FPGA as it is needed. It will appear as an adjunct to or an extension of the microprocessor.

ARC International was a pioneer in configurable microprocessors and it is now collaborating with Toshiba. ARC software enabled engineers to extend the microprocessor's instruction set to suit a particular application. It's a step in the right direction, but it is still design-time configuration.

The needed further step is to make the chip generic in manufacture (to amortize production cost across a range of applications) and configurable in the field. Microprocessors are generic in manufacture, but cannot be customized in the field. Microprocessors from ARC and **Tensilica**, which can be customized at design time, can achieve dramatic improvements in application performance. But design-time configuration fragments manufacturing scale. **Stretch, Inc.** makes generic Tensilica-based microprocessors with on-chip programmable logic. Generic in manufacture and customized in the field by programmers using Stretch's development software, these chips have potential to help bridge the gap between ASIC designs and microprocessor-based designs, while remaining accessible to programmers with no logic design experience.

Consequences of universal access

By universal access, I mean wireless access to information anywhere, any time, and by anyone.

Mobile devices. Every mobile device will have at least one radio and many will have several to facilitate transparent relay in ad hoc networks or to simultaneously communicate with disparate services. For example, billions of cell phones and other mobile devices will have GPS receivers and complementary inertial navigation systems to offer seamless position information anywhere.

Base stations. Ubiquitous networks mean proliferation of base stations that bridge mobile devices to the fiber network. More millions of radios. These will be agile radios, able to adapt to the requirements of transient mobile devices and to move around the spectrum to avoid interferers. They'll have smart antennas too, able to focus both transmission and reception.

Automobiles will have universal navigation aids with intelligent, engaging verbal interfaces. I should be able to say "Help me get to the airport. Avoid heavy traffic and construction and don't put me on any toll roads. I have thirty minutes." It should know which airport I'm talking about and it should know the best routes. It should be able to get real-time traffic information and integrate the information to get me there in time. If it cannot do that, it should say so and tell me the expected driving time.

Perhaps the car will take over the driving too. DARPA's 2005 challenge made autonomous driving look decades from solution. A year later, five autonomous vehicles successfully navigated the 131.6 mile course, with the three fastest cars averaging almost 20 mph on the rugged cross-country trek. The 2007 DARPA challenge is an urban course. Arrive in an unfamiliar city on a rainy night and you'll appreciate a car that knows the local area.

The hundreds of sensors aboard the auto will be on one of the vehicle's private and secure networks, eliminating miles of point-to-point wiring that today's cars lug everywhere.

Autos are due for an even more fundamental transition than just the proliferation of driver-aid electronics. For a hundred years, the drive train and braking systems have been mechanical. The inefficient mechanical drive train loses power between the engine and the wheel. All of the energy of braking is lost to heat. Motor/generators

at each wheel will make the vehicle lighter and more efficient. DC wheel motors produce infinite torque on start and, as generators, capture energy while braking.

Electric cars aren't only for fuel-conscious consumers who don't care about performance. See, for example, **Wrightspeed**'s X1 street-legal prototype car at *www.wrightspeed.com*. One design objective for the X1 was acceleration from 0 to 60 mph in 3.0 seconds, faster than internal-combustion sports cars, while still achieving efficiency equivalent to more than 170 mpg.

Once the network is everywhere, *sensors* will proliferate. The nuts, bolts, and cables holding bridges together can report their condition. Windows can summon the washer-bot. I've been accumulating stuff for decades and can no longer keep track of what I have or where it all is. With RFID tags, a network of sensors, and a mobile computer, I would be able to find that impact wrench I bought ten years ago and stored somewhere in the shop. Once all my stuff is tagged, my mobile device can tell me the next time I visit my neighbor Joe that the chainsaw he borrowed from me six months ago is still in his garage.

Consequences of emerging economies

China, India, and the rest of the emerging economies represent half of earth's population. That's *twice* the number of people in the developed world. Servicing demands of emerging economies will probably triple the size of the semiconductor market. Hundreds of millions of refrigerators, washers, microwave ovens, and other appliances mean billions and billions of electronic components at the lowest possible cost and with just enough performance to do the job. For the starter set of home appliances, power-efficiency isn't a consideration; it doesn't matter if the brain of a 1500-watt microwave oven or clothes dryer is a tenth of a watt or ten watts.

Unit volumes in the billions at the lowest possible cost and with performance that is just adequate means generic components—microprocessors, memories, peripheral chips—built of *value transistors*, the lowest-cost transistors that will do the job. For home appliances, that means almost any semiconductor process for which process development and plant cost have already been amortized to zero. Cost of production is the cost to run the foundry.

This means increased demand for silicon, for semiconductor processing equipment, and for semiconductor fabrication plants (fabs). This isn't demand for leading-edge equipment and state-of-the-art fabs; it's demand for used equipment and for new equipment designed not for leading-edge processes but for cost-effective processing. Semiconductor equipment makers will have to adjust from a long tradition of building high-cost, leading-edge equipment to building low-cost, modular equipment with standard interfaces. That's too much change for most equipment makers, so I expect new makers of value equipment to emerge and take that market. It's an opportunity for flexible second-tier vendors to make headway against the market leaders.

What it all means

To visualize market effects, I use a "Zeroes Model": zero cost, zero delay, zero power, and zero volume. These overlapping characteristics define the primary attributes of a product. For "zero cost," the designer's objective is to achieve the lowest product cost. For "zero delay," the designer's objective is to maximum performance—zero delay from request to result. For "zero power," the designer's objective is minimum power. "Zero volume" designs are show off productions for bragging rights or technology demos with an expected market of zero units.

Competing in the *zero-cost* segment dominant in emerging economies are consumer products, particularly appliances, where saving a nickel in the bill of materials can translate into millions of units in additional sales.

In the *zero power* segment are mobile devices, where the ultimate goal is to run continuously on scavenged power. A mobile device for the consumer market is in the overlap between zero power and zero cost. If the consumer's device also has high computational requirements, such as processing for digital media or cellular telephony, then its computational requirements place it in the *zero-delay* segment's overlap with zero power and zero cost. The overlap of zero delay, zero power, and zero cost emphasizes cost-performance per watt.

From 1980 until 2000, the elephant in the semiconductor business was the PC, which consumed as much as 50 percent of the dollar value of semiconductor output. Since 2000, the industry has been moving from the PC's cost-performance design to the mobile device's cost-performance-per-watt design.

For more than four decades, shrinking the transistor has been the recipe for semiconductor progress. A smaller transistor runs faster, it uses less power, and it is cheaper to make. A double whammy is about to hit this recipe for progress.

Applications for emerging economies want the cheapest transistors they can get, value transistors. Mobile devices, which are mostly idle, also don't want smaller transistors, which use less power when they're switching but leak more when they're idle. The search is on for better transistor geometries that aren't as leaky.

One alternative to shrinking transistors, poised for near-term commercial success, is *wafer stacking*. This process stacks two to more than ten wafers together with hundreds of thousands of vertical interconnects between chips. Stacking wafers dramatically reduces wire lengths between components. Resistance and capacitance of wiring is reduced, so drivers and receivers are smaller.

Circuits shrink, get faster, and use less power. There's less noise, less crosstalk, and less power in radiated losses. Analog wafers can be stacked with digital-logic wafers and with memory wafers. This permits independent redesign or upgrade; redesign of the digital logic would no longer require redesign of analog subsystems.

Wafer stacking also improves yields. Stack an extra processor wafer and an extra memory wafer or two and the probability of getting a good processor and memory rises. For example, stacking two 90 percent yield processors brings the expectation for at least one good processor to 99 percent, with similar gains from adding extra memory wafers.

Startup **Tezzaron Semiconductor** (formerly Tachyon Semiconductor) is a pioneer in wafer stacking. Other startups, such as **Tru-Si Technologies**, **ZyCube**, **Cubic Wafer** (formerly Xanoptix) and large companies, such as IBM, **Intel** (INTC), **Samsung**, and Toshiba, are experimenting too.

Nature and nanotechnology

One legacy of leading edge semiconductor processing equipment and instruments is to enable engineers to see and shuffle atoms and molecules. Merging the world of transistors with the worlds of biology and molecular chemistry, engineers and scientists are beginning to harness Nature's inventions such as the engines of cells, bacteria, and viruses. Natural systems capture energy, sense and manipulate surroundings, fail gracefully, and adapt in ingenious ways. Biological mechanisms will aid self-assembly of complex systems.

For the past few years, nanotechnology, nanomaterials, quantum dots, buckyballs, carbon nanotubes, and nanowires have been front-page stories. While nanoparticles have made their way into bulk materials such as sunscreens and textiles, those applications don't require a microchip's precise organization of millions or billions of carbon nanotubes. Demonstrating a carbon nanotube transistor, a nano-scale memory cell, or a nanowire in a lab is no guarantee that commercial application will follow.

More practical are MEMS (microelectronic machines), which are miniature systems combining electrical and mechanical parts. They will be the sensors and actuators in tomorrow's mobile devices. There are already some successes, such as airbag accelerometers, inkjet nozzles, automotive exhaust sensors, and gyroscopes from companies such as **Analog Devices** (ADI), **Freescale** (FSL), and **Texas Instruments** (TXN). But the journey from lab to commercialization has been difficult as engineers struggle with new problems (e.g., stiction, mechanical release, deep etching) and with new techniques and materials. Also, there are no generic components that can achieve huge unit volumes like microprocessors and memories. MEMS designs tend to be application specific. Application-specific designs are expensive and have limited production volumes, further increasing costs.

Nutshell forecast

Cost-performance per watt and time to market will dominate design for the consumer systems that make up the bulk of semiconductor demand. The semiconductor market will continue to grow in the range of 10 to 15 percent per year as electronics invades everything from carpeting to car bumpers. Demand from emerging economies will spur growth in the traditional semiconductor market. At least one Holy-Grail memory cell, with the speed of SRAM, the density of DRAM, and the non-volatility of flash memory, will emerge and gain commercial importance within five years. Wafer stacking and other techniques will displace transistor-shrinking as the means to continue historical performance improvements and price declines in semiconductors. Self-identifying serial interfaces will become the way to connect disparate subsystems in single-chip designs, simplifying integration of functional blocks. Programmable-logic derivatives will become pervasive in the form of systems that house physical circuits in the way that today's memory devices house programs. These programmable-logic chips will be generic at manufacture and customized in the field.

– Nick Tredennick and Brion Shimamoto
June 28, 2006

Got Questions?

Visit our subscriber-only discussion forum, the Telecosm Lounge, with George Gilder and Nick Tredennick, on www.gildertech.com

GILDER TECHNOLOGY REPORT PUBLISHED BY GILDER PUBLISHING, LLC AND FORBES INC. • Copyright ©2006, by Gilder Publishing, LLC

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

EDITOR	TECHNOLOGY ANALYST	MANAGING EDITOR	CONTRIBUTING EDITORS	ART DIRECTION	SUBSCRIPTION DIRECTOR	CUSTOMER SERVICE
George Gilder	Charlie Burger	Mary Collins	Brion Shimamoto Nick Tredennick	Thomas McCarthy	Rosaline Fernandes	Sandy Fleischmann

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