GILDER
TECHNOLOGY
REPORT

April 2007 / Vol. XII No. 4

What's Next?

To get to the cell phone of the future, with local intelligence, extended battery life, sensors, better network connectivity, and inertial navigation, the semiconductor industry will have to embrace MEMS and semiconductor wafer stacking.

Inside:

We meet for breakfast at 0600 on Thursdays. I was gloating "Intel announced an ovonic memory chip." Keith Diefendorff knew, of course, that I had predicted a few years ago ("Goldilocks and the three memory chips," *Dynamic Silicon*, May, 2002) that ovonic unified memory (OUM) was the likely winner in the non-volatile memory sweepstakes. His immediate response was: "What's next?" His question surprised me (I was expecting something on non-volatile memory) and it set me on the path to this missive.

I'll enumerate a few obvious trends. Next, I'll speculate about their consequences. From there, I'll specify the technology that supports the changes. Then it gets more difficult as I bridge from the technology forecast to companies and to specific investments. I'm hoping to run out of print space before I get to the difficult sections.

Obvious trends

Here are the trends that will influence semiconductor development, in no particular order: rapid growth in handheld electronics, ubiquitous Internet use, wireless access everywhere, the shift from business products to consumer products as the major market for electronic components, the change from analog to digital, and the growth of electronic transactions. No big deal here, we all see these trends and more.

Consequences

The consequences of these trends—what constitutes "better"—for the semiconductor industry include: miniaturization, power economy, antenna improvements, more integration, cost reduction, and lots of sensors.

This seems like business as usual for the semiconductor industry, but it isn't. The shift from business products to consumer products as

(CONTINUED ON PAGE 3)

FEATURED COMPANY: Energy Conversion Devices (ENER)

Don't Convert to ENER—Yet

Energy Conversion Devices (ENER) had been working on its chalcogenide-based phase-change memory technology, called ovonic unified memory (OUM), for more than two decades when in 1999 pivotal **Micron** (MU) executives Tyler Lowrey and Ward Parkinson co-founded **Ovonyx** in order to commercialize Energy's proprietary memory chips. Veterans of the semiconductor memory industry, Parkinson co-founded Micron in 1978 and served as its first chairman and CEO, while Lowrey led operations and research for the ascendant memory maker over a period of 14 years.

When Ovonyx was formed, Energy exclusively licensed its memory technology to the venture, which it jointly owns with Lowrey, **Intel** (INT), **BAE Systems**, and **STMicroelectronics** (STM). In return for its one-third ownership of Ovonyx, Energy has invested $1.5 million in the company while reaping losses of $1.3 million, including trickles of revenue from services provided to the joint venture and from the 0.5 percent of annual gross revenue that Energy is entitled to receive. (Ovonyx is paid license and service fees as well as royalties.)

But after years in the desert, OUM may be ready to blossom with the promise of reading and writing data 500 times faster than today's flash memory along with the prospect of extending the battery life of portable devices. Licensee Intel expects to sample 90-nanometer, 128-megabit OUM-based memory before summer, with mass production to follow by the end of the year. The chip can survive over 100 million read-write cycles and store data for longer than a decade.

As part of the greening of OUM, Ovonyx expects to secure more licensees. For instance, late last year the venture granted a new license to **Qimonda** (QI), formerly the memory operation of ***Infineon*** (IFX), to commercialize OUM through a joint research program with **IBM** (IBM) and **Macronix** (MXIC) of Taiwan. The companies demonstrated a prototype phase-change memory device with dimensions of just 20 nm, a geometry not expected to be fabricated on a mass scale for another eight years.

Can you cash in on OUM?

Unfortunately for investors, OUM has a dark side—Energy Conversion Devices itself. Stanford and Iris Ovshinsky founded the company in 1960 in order to realize their dream of mass producing photovoltaic (PV) materials, which convert sunlight to electricity. Since then, Energy has grown like an octopus, extending its tentacles into all sorts of futuristic and dead-end technologies, such as Ovonic cognitive computers (including Ovonic quantum control devices), hydrogen storage for automobiles, and fuel cells.

Still at the heart of the company is the volume manufacture of thin sheets of PV materials, which spool at every increasing speeds out of machines that span the length of a football field. The stainless steel rolls, each about a mile and a half long, currently take about two and a half days to manufacture using lighter, cheaper, more flexible materials than competitors. By perfecting a process that requires only minute amounts of the expensive silicon crystals that plague most of today's PV products, the Ovshinskys improved more than the processing; they also created material that absorbs a wider spectrum of solar radiation and that is easier and cheaper to transport and install.

United Solar Ovonic, Energy's wholly-owned subsidiary for the manufacture and sale of PV modules and already the largest domestic producer of PV materials, has been rapidly expanding manufacturing capacity as it continues to build new plants and add capacity to older ones. The company has doubled capacity since last fall and plans to double it again by the end of this year. In addition, United Solar just agreed to a joint venture (75 percent owned by United) with **Tianjin Jinneng Investment Company** of China to assemble PV modules from the solar cells manufactured in Michigan.

Energy also manufactures positive electrode nickel hydroxide battery materials and licenses NiMH (nickel metal hydride) rechargeable battery technology. Through its 91 percent owned subsidiary, Ovonic Battery Company, Energy has formed Cobasys, a strategic alliance with partner **Chevron** (CVX), which currently funds Energy's portion of the venture. The only US-based supplier of NiMH battery systems for hybrid electric vehicles, Cobasys was selected by **General Motors** (GM) to provide systems for the 2007 Saturn Aura Green Line Hybrid mid-size sedan, available this spring. GM has also awarded Cobasys a contract to develop and test lithium-ion battery system technology for the GM plug-in hybrid electric vehicle program. Partner A123Systems will supply the lithium ion batteries based on its proprietary nanophosphate technology, and Cobasys will work on integrating and manufacturing the lighter-weight system appropriate for smaller vehicles.

Invest in your peace of mind instead

United Solar Ovonic, which continues to account for nearly all product sales, contributed to 84 percent of Energy's revenue in the December quarter, generating operating income of $1.15 million (down from $1.76 million during the year-earlier quarter). The remainder of Energy reported combined losses of $8.5 million, resulting in total operating loss of $7.9 million, a bit worse than the loss of $7.1 million the previous December. During the same period, revenue fell to $22.9 million from $24.3 million.

Haunting Energy recently have been cost overruns associated with United Solar's build-out along production problems at the newest PV facility and falling demand from Germany, which accounts for about half of United Solar's PV sales. Worse, Cobasys has lost $225 million since its inception in 2001. Most recently, the subsidiary lost $21.7 million during the December quarter, double the year-ago loss. Energy records only small amounts of revenue for services performed on behalf of Cobasys, and now faces the dismal prospect that partner Chevron may cease funding the venture this year.

If you Google NiMH, the first hit is the National Institute of Mental Health, the agency that stands ready to aid all those who invest in ENER today in an attempt to profit from the possible success of Ovonic Unified Memory tomorrow. The company has chronically lacked focus and direction, resulting in a feat perhaps no other corporation can boast of—after more than forty straight years of losses, the company continues on in a financial version of eternal life. Though Iris died last summer and Stanford recently stepped down from the posts of president and chairman, the corporate culture lingers on; despite the recent successes of OUM and potentially increasing leverage of PV production, management appears to be turning more of its attention to its fanciful hydrogen production and storage projects.

If you Google NiMH, the first hit is the National Institute of

Mental Health, the agency that stands ready to aid all those who invest in ENER today in an attempt to profit from the possible success of Ovonic Unified Memory tomorrow. The company has chronically lacked focus and direction, resulting in a feat perhaps no other corporation can boast of—after more than forty straight years of losses, the company continues on in a financial version of eternal life. Though Iris died last summer and Stanford recently stepped down from the posts of president and chairman, the corporate culture lingers on; despite the recent successes of OUM and potentially increasing leverage of PV production, management appears to be turning more of its attention to its fanciful hydrogen production and storage projects.

At around $36, Energy's enterprise value is a remarkably rich $1.14 billion or 11 times annual revenue of about $100 million. If you want to invest in OUM, our advice is to follow the technology and wait until you are reasonably certain ENER will prosper because of it—a scenario that will probably unfold more by accident than by plan.

— Charlie Burger

(CONTINUED FROM PAGE 1)

the major market for electronic components, in particular, signals a change in business models.

Game developers are emblematic of the problems that companies have in producing consumer-oriented products (as opposed to business-oriented products). Product cycles are short, typically less than one year. Developers of game platforms exhibit no component loyalty. The microprocessor at the heart of this year's game platform has no seeming advantage in competition for a place in the next system from the same manufacturer. Product volumes are high, but component margins are low because platforms may be offered as below-cost enticements to game buyers.

Short product cycles and the absence of brand loyalty are hallmarks of the HDTV makers and the same probably goes for cell phone makers. There's a shift in semiconductor unit-volumes, from office machines and business products with lengthy product cycles and stickiness to designed-in components to consumer devices with product cycles of a few months and no component loyalty.

For businesses this means two things. First, it is difficult to build customer loyalty. But your business-rivals' customers will be more accessible. Second, your products need to be synchronized with the product cycles of your major customers. If your major customer produces a new game platform every eight months, your new products are guaranteed to miss more opportunities if they are on a one-year cycle than if they also run on an eight-month cycle that is synchronized with, but ahead of your major customers.

The Future

This year's luxury is next year's necessity. Automobile subsystems that seemed exotic a few years ago, navigation systems, entertainment displays, heads-up displays, engine and environment sensors, and rear-looking cameras, become more common and will soon be regular features. Automobiles will acquire massive local intelligence: location awareness, situational awareness, and wireless connectivity.

If our automobiles can do this, can our personal devices be far behind? The limited local intelligence, primitive location awareness, and wireless connectivity that handheld devices already have will improve and they will soon acquire situational awareness.

I say current location awareness is primitive because it just isn't good enough for coming applications. GPS works only with a mostly unobstructed view of the sky; it's not good in buildings or tunnels. Triangulation by base stations is OK, but lacks precision, altitude information, and it doesn't work everywhere. It's not good enough for my cell phone to know I am at a shopping mall, it should know where I am in the mall and it should help me locate the store I'm looking for. It's not good enough for my cell phone to know that I am in a particular building; it should know the floor I am on and the floor I am headed to so it can help me arrive on time. Location awareness isn't just knowing latitude and longitude with precision; it's knowing what is in the vicinity and why that might be important to me.

Sitting at my desk, I've gotten used to having at my fingertips any information I want. Next, I want that information in my car and from my cell phone. Eventually, the second-hand information I have access to now won't be good enough either. It won't be sufficient to check conditions at the ballpark or zoo on weather.com or to depend on broadcast traffic reports, I'll want access to the video from

Advanced Micro Devices	(AMD)
Altera	(ALTR)
Anadigics	(ANAD)
Broadcom	(BRCM)
Cepheid	(CPHD)
Corning	(GLW)
Energy Conversion Devices	(ENER)
Equinix	(EQIX)
EZchip	(LNOP)
Finisar	(FNSR)
Flextronics	(FLEX)
FormFactor	(FORM)
Hittite Microwave	(HITT)
Ikanos	(IKAN)
Micron	(MU)
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)

Microvision (MVIS)

PARADIGM PLAY: HEAD-UP TELEPUTER DISPLAYS

APRIL 11: 4.09; 52-WEEK RANGE: 1.16 – 4.35; MARKET CAP: 176.66M

Would you pay an extra hundred dollars for a cell phone with an embedded projector? Alexander Tokman is betting that about 3%–5% percent of those who purchase "feature rich" phones will jump at the chance, meaning unit sales of mobile phones containing his PicoP microprojector could reach 24m–40m annually based on last year's figures. The feisty CEO of Microvision also hopes to see PicoP in other handset devices such MP3 and DVD players and to sell it as a standalone accessory the size of an iPod. Imagine carrying in your pocket a projector that can expand your cell phone's 2" screen into a 30" or 100" screen of bright, crisp images (even when projected onto distorted surfaces) to watch mobile television and video or to browse the internet.

Tokman imagines millions of consumers doing just that before the end of next year. But first he has to overcome the manufacturing hurdles still blocking his tiny display engine. (We described these in detail in the January GTR along with Microvision's world-beating display technology.) Tokman exudes Lombardi-like confidence that he's about to execute a game-winning touchdown. His offense includes the likes of Corning, Novalux, and Osram working feverishly on green lasers specifically for PicoP, two major high-volume manufacturing partners gearing up for the big ramp, and 75 in-house engineers (over half of his total staff).

During the consumer electronics show in January, Tokman met privately with 31 manufacturers and expects to reach development agreements with several this year. These same manufacturers tell Tokman that he has virtually no competition based on size, cost, power, and image brightness and contrast. As the PicoP ascends rapidly into handsets throughout 2009, Tokman expects his technology to begin appearing in luxury cars as reprogrammable, instrument-cluster displays. The added benefit for manufacturers of automobile modules is elimination of the expensive manual process required to install the inferior display systems offered by his competitors.

Over the past year Tokman cut Microvision's headcount by almost a quarter while increasing its engineering force by a third. Thanks to this aggressive streamlining and refocusing on the most promising consumer applications, he has enough cash to keep his company running for about 10 months. But that's still over half a year before Tokman expects to begin reaping his first serious revenue, and thus he will be forced back to the capital markets this year.

However, even if he offers his entire $35m shelf (at the recent stock price of $4 that would dilute shares outstanding by a fifth to 52m), Microvision's resulting market cap of $208m is still dirt cheap if Tokman wins his Super Bowl of tens of millions of unit sales.

Sigma Designs (SIGM)

PARADIGM PLAY: LIFE AFTER TV MEDIA PROCESSOR LEADER

APRIL 11: 26.14; 52-WEEK RANGE: 7.99 – 32.57; MARKET CAP: 595.84M

The now familiar Sigma growth story continues unrestrained as its ascendant media processors flood the swelling digital-media markets. Twelve telcos worldwide are currently deploying IPTV services with Sigma inside their set-top boxes, and that number is expected to double over the next year. Meanwhile, the Blu-ray format appears to be overtaking the new generation of high-density video disks, where Sony, Pioneer, Panasonic, and Sharp already rely on Sigma. Look for Sigma's Blu-ray revenues to begin ramping strongly during the second half of this year. Sigma processors are also seeing early sales into HDTVs, portable media devices, and digital adaptors for home media centers.

Management believes they have captured three-quarters of the IPTV set-top box market, which they expect to double this year. If Sigma's IPTV revenues correspondingly double, they could reach $47.4m during next year's fourth quarter ending January. Though unit prices may ease as volumes ramp, we believe this estimate is likely conservative for two reasons. First, Sigma sees signs, based on its customers activities, that these growth estimates may be low. Second, the company has a good shot at increasing its market share as the world of IPTV moves toward the multiple codec solution that Sigma has mastered and that rivals continue to struggle with. Sigma remains the only supplier to the Microsoft TV platform which has begun to ramp in Europe and North America and which could rise to a third of IPTV unit volume by the end of this year.

Of greater concern to stockholders than competitors or markets is Sigma's inquiry into its past options-granting practices. Spurred on by a new CFO and fresh audit firm, we can only hope the internal inquisition is fast concluding as the clock ticks toward the NASDAQ's 4 May reporting deadline. Until then, we must guess about earnings. Management did say that its target gross margin of 50% has been reached as cost reduction efforts and quantity discounts from increasing wafer volumes are offsetting volume price discounts. But the company refused to defend its prior expectation that operating margin would eventually hit 20%. Sigma may well be feeling the cost burden of the inquiry while recognizing the need for greater internal oversight going forward.

If Sigma's non-IPTV revenues rise a modest 25% this year, they could easily report a $57m quarter next January for a $228m annual run-rate with the prospect of continued ascent into the next year. If we grant Sigma an operating margin of 18% and a growth valuation of 30x earnings, then the stock price would be scraping $41 next March based on 30m shares. Not long thereafter, however, income taxes should begin biting Sigma, a looming menace even bigger than the options inquiry. Thus, while we think the stock may eventually reach the 50s, it may be quite some time before you see another double at the current price of $26.

Watch this one closely; the swells—and attendant swoons—typical of high growth stocks may present more buying opportunities over the coming months while we await clarity on the financials, including margins, earnings, and taxes.

— Charlie Burger

the cameras at the ballpark (what's the weather?), the parking lots (is it too crowded?), and the roads (is traffic flowing freely?). My cell phone should know the weather where I am and the weather where I'm going (it has, after all, access to my appointment calendar) and it should alert me to differences.

In addition to all this, my cell phone is heavy and bulky and its battery doesn't last. The next time you think your cell phone is small enough and light enough, throw onto the scale all the stuff you pack when you travel with it.

So how will the semiconductor industry get from today's cell phone to the cell phone of the future? By adding functions: local intelligence, battery life, sensors, better network connectivity, and inertial navigation. To get there, the semiconductor industry will have to embrace microelectromechanical systems (MEMS) for the sensors, gyroscopes, accelerometers, and RF components such as oscillators. And it will have to embrace semiconductor wafer stacking.

Microelectromechanical systems (MEMS)

Microelectromechanical systems are chip-scale systems that combine mechanical subsystems with integrated electronics. Think of a regular electronic integrated circuit that includes a few moving parts. An accelerometer, for example, might combine a diving-board-like feature with sensitive electronics that measure the changes in capacitance as the diving board moves closer to or farther from capacitive plates as a result of motion of the system that contains the chip.

One key advantage for MEMS chips is that they benefit from the same batch manufacturing processes that have made integrated circuits so successful. It has taken a long time, however, to refine manufacturing techniques to deal with substantial differences between batch fabrication of mechanical systems, which have three-dimensional structures and moving parts, and fabrication of planar electronic integrated circuits. The MEMS makers are finally getting to components with costs that are competitive with discrete components, but that have better performance and higher reliability.

Four application areas drive the ascendance of MEMS: sensors and actuators, RF components, biomedical applications, and optical networking. MEMS chips were a hot topic for a while, until nanotechnology stole the media attention. Despite waning public attention, the MEMS market has been developing nicely. According to Yole Developpement (http://www.yole.fr), the worldwide market for MEMS components was $5.3 billion in 2005 and is expected to grow to $9.8 billion by 2010. Consumer applications in cell phones (microphones, resonators), in television sets (optical devices), and in game systems (gyroscopes, accelerometers) drive this compound annual growth rate of 13 percent.

ADI is the granddaddy of MEMS accelerometer makers. Other long-time players include Bosch, Freescale, Infineon, STMicroelectronics, and TI.

MEMS components are the right thing at the right time for the burgeoning mobile-device market. MEMS components are the same size as semiconductor chips that are the heart of mobile devices. They offer increased sensitivity and increased reliability and, because they are integrated with the electronics, they enable new design techniques. MEMS sensors and actuators can be designed to fit the application. They're not so large that they waste power and board space and they're not too small to be effective. For example, MEMS actuators in the form of speakers and microphones, are large enough to interact with humans, but small enough to be treated as chips for assembly purposes. A sensor that's too small cannot collect enough samples to be representative; a sensor that's too large collects samples that are larger than necessary and it uses more power.

Akustica, a fabless semiconductor company, makes digital MEMS microphones. **Knowles Electronics** and **Sonion** make analog MEMS microphones. **SiTime**, a fabless semiconductor startup, makes MEMS resonators and oscillators for RF clocking and timing. SiTime's MEMS resonators can replace quartz crystals in radios. Startups **Discera** and **Silicon Clocks** also make MEMS oscillator chips.

Gyroscopes and accelerometers will probably be the next high-volume MEMS components to be ingested by the cell phone after MEMS microphones and resonators. Accelerometers sense changes in linear motion; gyroscopes sense the rate of rotation about an axis.

Analog Devices (ADI) is, of course, the granddaddy of MEMS accelerometer makers, listing thousand-unit prices as low as $3.50 for two-axis accelerometers on its web site. Analog Devices also has a line of gyroscopes, listing for as little as $22 in thousand-unit quantities. The high-volume price for a three-axis accelerometer is about $2. Analog Devices and **STMicroelectronics** (STM) have shipped millions of three-axis MEMS accelerometers in **Nintendo** Wii game systems.

Because they are flat, today's chips aren't efficient. Imagine the area inefficiency if Hong Kong or New York were to abandon skyscrapers. It's easy enough to do: that's Los Angeles.

Analog Devices got its start in MEMS by designing accelerometers for air bags. Other long-time players in the MEMS market include **Bosch**, **Freescale** (FSL), **Infineon** (IFX), STMicroelectronics, and **Texas Instruments** (TXN). **Hewlett-Packard** (HPQ) has been in the MEMS business for some time making inkjet nozzles. The sophisticated chips that drive its inkjet cartridges incorporate half-a-million transistors on the same chip with 3900 ink nozzles. HP recently announced that it is working on MEMS accelerometers and gyroscopes for consumer applications such as cell phones.

Late last year, **InvenSense**, a fabless semiconductor startup, announced a MEMS gyroscope with a projected price of $2 per unit. It has two products, the IDG-300 and the IDG-1000, in production now. The gyroscope's applications are image stabilization, navigation, and gesture-based applications (tilt the device to scroll through a menu, for example).

The next time you think your cell phone is small enough and light enough, throw onto the scale all the stuff you pack when you travel with it.

For investment advice, it's the usual problem. The large companies are so big that a healthy MEMS market may not have much effect on the stock price compared to other things the company does. The small, innovative companies are startups that offer no simple means for investing.

Wafer stacking

Because they are flat, today's chips aren't efficient. Imagine the area inefficiency if Hong Kong or New York were to abandon skyscrapers. It's easy enough to do: that's Los Angeles. Today's chips are Los Angeles. Shrinking transistors helps, just as shrinking houses would reduce the sprawl of Los Angeles.

The problem is the distances that signals have to cover cost time and energy. Greater distances mean signals take longer to travel and it takes more electrons to create the right voltage on the wire.

Shortening interconnecting wires improves everything. Shorter wires fill and empty faster and they don't require as many electrons to reach switching voltages. In addition, the transistors can be smaller because they don't have to

When a transistor switches on, electrons flow through the conductors connecting that transistor to its neighbors. Think of it as an irrigation system; when a sluice gate is open, water flows into the canal. CMOS is complementary, so as one canal is filling another is emptying. CMOS works by detecting the levels in the canals, not by measuring flow in the canal. The greater the distance between sluice gates, the longer it takes for the water to reach the next sluice gate or to drain from a full canal. Also, the longer the canal, the longer it takes to fill or empty simply because the volume of a long canal is greater than the volume of a shorter one. In CMOS the electrons are like the water and the connecting wires between transistors are like the canals; the transistors are, of course, the sluice gates.

push as many electrons around. Smaller transistors use less power. The idea of wafer stacking is to move from flat circuits and wiring to three-dimensional circuits with vertical wiring.

A typical integrated circuit (chip) contains kilometers of wires. Signal paths can be tens of millimeters long. In modern chips, the interconnect wiring dominates chip area and power requirements and it limits performance. Quarter the chip and stack the pieces. The vertical distance to connect two quadrants is a hundredth of a millimeter (10 microns)—a distance so short it can be neglected relative to planar distances measured in millimeters. Stacking the wafers halves communication distances (a 100mm × 100mm chip becomes a stack of four 50mm × 50mm chips). Chips in these wafer stacks can have hundreds of thousands of these short vertical connections, called through-silicon vias (TSVs), without significantly adding to each chip's area. TSVs can be as small as a micron in diameter and are typically 10 microns long.

Since wafer stacking shortens interconnects, Samsung finds that its wafer-stacked 16-Gb flash memory is about 30% faster.

Intel's experimental 80-core microprocessor uses TSVs to communicate with its memory. This fast, close TSV connection between the microprocessor and its memory eliminates perhaps the most performance-limiting bottleneck in today's microprocessor systems (communicating through a memory controller and over standard chip-to-chip interfaces). Because of the wide, fast data connections enabled by TSVs, transfer rates between the processor and memory can reach a terabyte/second. Intel CTO Justin Rattner has said that Intel is working on TSVs for a variety of products.

Stacking enables mixing analog circuits, sensors, logic, software (in ROM, for example), and memory. Separating the analog circuits from the digital circuits permits the digital circuits to progress with Moore's law without forcing redesign of the analog circuits (analog circuits do not scale as easily as digital circuits). Separating memory from logic has similar advantages. Logic transistors need speed and generally care somewhat less about leakage; memory transistors want low leakage even if there's a speed penalty. The ability to mix chip types in a stack means that each chip can be optimized for its purpose. That is, there need be no compromise in a memory semiconductor process to allow it to accommodate logic circuits.

The yield of a stack, which would be the product of the yields for each individual chip plus an allowance for the yield of the stacking process, would seem to decrease with each added layer. Stacking just seven chips, each yielding at 90 percent, drops the result below 50 percent. That sounds dismal, but stacking also facilitates redundancy, which can drive the stack's yield above the yield of even a single chip. Stack an extra memory chip and design the circuit to self-test and self-organize on startup. Stack nine chips in a byte memory. When the chip initializes, it self-organizes into a byte memory, potentially ignoring a bad chip. The yield for this stack is greater than the yield for an individual chip.

Samsung is experimenting with wafer stacking. Since wafer stacking shortens interconnects, Samsung finds that its wafer-stacked 16-Gb flash memory is about 30 percent faster. Samsung expects to ship wafer-stacked chips in NAND–based memory modules early this year. Samsung has said that it will also do system-in-package chips and high-capacity DRAM. Samsung isn't alone. A partnership among **NEC Electronics** (NIPNY), **Elpida Memory**, and **Oki Electric** has announced stacked flash chips. **Micron** (MU) is doing wafer-level packaging that uses TSVs. Freescale is working on TSV interconnection. IBM (IBM) is developing 3-D chips with a virtually unlimited number of interconnections between its layers.

Tezzaron Semiconductor is a fabless semiconductor company with several wafer-stacked products. It has 3-D stacked SRAM, 3-D stacked DRAM, and 3-D stacked SRAM with an 8051 microcontroller. Tezzaron, which was featured in my article of five years ago (back then it was Tachyon Semiconductor), has been a pioneer in the development of wafer stacking. Bob Patti, Tezzaron's CTO then and now, knows why wafer stacking has taken a long time to develop. It isn't because it is particularly difficult technically; it has more to do with specialization of job functions.

As John Trezza of the pioneering **Cubic Wafer** of Austin (formerly Xanoptix) discovered, (see *Gilder Technology Report*, January 2006) design engineers and manufacturing engineers live in separate worlds. When Tezzaron and Glenn Leedy's **Elm Technology**, a holder of fundamental patents in wafer stacking, went to large semiconductor companies to talk about wafer stacking, they typically got shuttled to the manufacturing (process) engineers. Wafer stacking does look like it belongs to the process engineers. But wafer stacking requires participation by the design engineers, who not only weren't meeting with Tezzaron and Elm, but had no incentive to deviate from techniques that were working perfectly well for them (shrinking transistors). The processor-memory bottleneck, shrinking system sizes, and the search for lower power consumption are finally changing designer incentives. Even when the designers are on board with wafer stacking, there'll be a delay as the electronic design automation (EDA) software companies revise their products to support the design for stacked wafers.

Equipment vendors **Alcatel** (ALA), **EV Group**, **XSiL**, and our favorite **Semitool** (SMTL) formed the EMC-3D consortium with the goal of lowering the cost of the 3-D process below $200 per wafer. I discussed wafer stacking five years ago ("Life After Moore's Law," **Dynamic Silicon**, Jan 2002). That was early; if the equipment makers are thinking about it, it's probably about to take off.

There's more

Back to the beginning: What's next? My answer to Keith was: "MEMS, wafer stacking, and RFID." For MEMS, I covered a few examples in RF components and in sensors and actuators. There's much more, such as the devices from **Microvision** (MVIS) and Texas Instruments, in optical MEMS and in MEMS for biomedical applications (see page 4). I covered RF components and sensors and actuators because they benefit first and most from the semiconductor market shift to consumer-products. The demand for wafer stacking is a consequence of the transition to mobile devices and the need for more performance at lower power. RFID belongs in the "what's next" answer, but demand for it is driven by different trends.

– Nick Tredennick, with Brion Shimamoto
April 11, 2007

GILDER TECHNOLOGY REPORT PUBLISHED BY GILDER PUBLISHING, LLC AND FORBES INC. • Copyright ©2007, by Gilder Publishing, LLC

291A MAIN STREET, GREAT BARRINGTON, MA 01230, TEL: (413) 644-2100, FAX: (413) 644-2122 • EMAIL: INFO@GILDER.COM

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

For subscription information telephone toll free: (877) 733-7876 • WWW.GILDERTECH.COM