GILDER TECHNOLOGY REPORT

September 2004 / Vol. IX No. 9

The Microprocessor Assault

Who will win the long-running see-saw battle between Altera and Xilinx?

Inside:

My neighbor Joe told me he was bidding on **eBay** (EBAY) for an all-terrain vehicle. I knew he wanted high performance, so I assumed it was a two-stroke engine. For decades it has been an axiom in the industry that two-stroke engines far excel four-stroke engines in performance. He saw my surprise when he told me he was looking for a four-stroke. "You don't understand," he said. "All the development money goes into four-strokes. They now outperform the two-strokes." The regulatory environment tilted research in favor of four-stroke engines and their efficiency and performance surged past the two-strokes. Times change.

The same thing is happening today in the semiconductor industry. For decades it has been axiomatic that custom logic and application-specific integrated circuits (ASICs) exceed microprocessors in raw performance for all specialized tasks. But for the last thirty years, semiconductor researchers have focused their efforts on microprocessors, steadily improving their design method, flexibility, performance, and power-efficiency. Innovation moves much faster in chips than in automobiles. Now the latest microprocessors are creating an upheaval in the industry that dwarfs the impact of four-stroke engines on the auto industry. In chips, it's as if they held an all-terrain auto race and someone entered a helicopter, with the same stunned shock, charges of cheating, and claims of conspiracy from the other drivers.

Initially the storm was confined to the esoteric world of microprocessor benchmarking—a microchip Formula One where engineers compete in running their devices through a suite of representative software. But soon the storm will hit the precincts of the nation's technology investors, and a portfolio near you.

At the center of the uproar are **Tensilica Corporation** and the Embedded Microprocessor Benchmark Consortium (EMBC, pronounced "embassy"). Tensilica is an intellectual property startup company that sells microprocessors that are configurable. A conventional microprocessor, such as the Pentium in your PC, has a fixed set of "instructions" or repertory of functions it can perform. Those instructions barely change over decades. But Tensilica offers a standard base microprocessor together with tools that allow hardware designers to add new custom instructions specially adapted to their applications. The chips hit the fan when Tensilica's configurable devices, specially adapted to the benchmark tests, blew away all the competition from such established microprocessor companies as **Texas Instruments** (TXN), **IBM** (IBM), **MIPs Technologies** (MIPS), **ARM** (ARMHY), **Freescale** (FSL) (previously Motorola semi) by a factor of three or more. How did it happen? Is it just cheating. Or is it a revolution?

CONTINUED ON PAGE 2

EMBC is just what its name implies—a nonprofit industry consortium, with about sixty companies, that publishes certified benchmark results for embedded microprocessors. EMBC's current benchmark suites include automotive/industrial, consumer, networking, office automation, and telecommunications, each with three to sixteen application programs. Using the results as a guide to microprocessor purchases are hardware systems designers. Difficult, contentious, and imprecise the benchmarking process may be, but it is much more reliable and revealing than merely comparing data sheets and marketing claims.

EMBC provides two types of benchmark scores: "out-of-the-box" scores and "full-fury" scores. For the out-of-the-box scores, the manufacturer must compile and run the benchmarks without changing the source programs, which are written in the prevailing high-level software language "C." If you can't change the code, it can't "see" or tap any changes on your chip. Explicitly barring any manipulation of the software to favor a particular micro's instructions, the out-of-the-box scores give the most pure and direct measure of a microprocessor's performance.

Five years ago, I wrote about the coming death of the DSP. It's closer today.

By contrast, for the full-fury scores, the manufacturer is free to optimize the source programs for their own microprocessor instruction set, rewriting the code to recognize new instructions, or even using a lower-level language called Assembler that hones the basic physical "machine language" of the chip. The full-fury scores show the performance enhancement that is possible with effort comparable to what the engineers will do for the final application in a real system when they adapt their software programs to run on a particular microprocessor.

Death of the DSP

The uproar first broke out in 2001 when Tensilica posted industry-leading full-fury benchmark scores. The Tensilica "base" processor was nothing special on the out of the box tests. But with additions to its instruction set and adjustments to the source code to exploit them, the full-fury Tensilica device improved on the base performance by a factor of 10 to 1,200. The average speedup for the telecommunications benchmark was almost 40. This should have persuaded the industry of the promise of configurable microprocessors. But it didn't. On the other hand, no helicopters are showing up in Formula One auto races either.

The plot thickened earlier this year when Tensilica posted similarly astounding out-of-the-box benchmark scores without the manipulations needed to map the benchmark programs onto the customized microprocessor. This time Tensilica's software alone was able to adapt its micro to the benchmarks in real time with no engineering intervention required. Taking advantage of "automagic" compiler technology, it translated the high level language of the source code on-the-fly into machine language that uses the new instructions. Tensilica's design software spits out a compiler that exploits custom instructions to speed execution even though it begins with unmodified source code. The speedup was three times the best out-of-the-box scores achieved by standard microprocessors, including TI's digital signal processors (DSPs) designed specifically for the application. Five years ago, I wrote about the coming death of the DSP. It's closer today.

What configurable microprocessors do is take the pivotal phases of microprocessor design away from microchip logic designers at places like TI and ARM and push them to software programmers at places like **Nokia** (NOK) and **Samsung**. That's huge because there are ten times as many programmers as hardware logic designers, and it's the programmers that shape the microprocessor-based systems that account for the billions of microprocessors that manufacturers ship every year. After all, no one buys a microprocessor simply to churn its instructions through benchmarks. Customers buy microprocessors as part of systems that execute software applications, which perform actual tasks in the real world.

Reconfiguring the industry

The relations between software and hardware have been reversed. In the past, computer manufacturers sold expensive computers and they gave away compilers and operating systems. Forty years of Moore's-law progress have driven the cost of hardware toward zero. The microprocessor is ubiquitous. Problem solving, once the domain of hardware designers, is now almost exclusively the domain of programmers.

No Moore's-law miracle, however, has hit the cost of programming. Applications, compilers, and operating systems—the fruits of programming labor—have become so expensive that changing the underlying hardware is usually out of the question. Desktop computers now are tied into the x86 microprocessor instruction set and cannot migrate away. Even embedded micros in digital cameras, cell phones, and automatic transmissions may have tens of thousands to millions of lines of programming. All of the programs use particular immutable instruction sets, requiring retention of the same microprocessor for subsequent product generations.

As the configurable processor advocates saw, this situation is ripe for change. It makes the microprocessor's creators merely guess at the computing resources and instructions that its users will need. It makes the programmers settle for whatever instructions are available. Then for the next generation, the beat goes on. Chip designers study software applications to find which instructions programmers are using and accelerate those instructions. Programmers pore through the manuals and use the instructions that the builders have accelerated. It's a viscous circle that fails to respond to what the programmers need when they need it.

Microprocessors are still mostly a take-it-or-leave-it deal. Large customers can ask the manufacturer for changes. But at

best these changes must defer to the manufacturer's two-plus-year development cycle and often are diluted by the design team and by the need to accommodate other suggestions. And when implemented, any improvements are available to everyone, including all your competitors.

Tensilica and its followers are reconfiguring this industry predicament. Here's the recipe for creating a Tensilica microprocessor. Take a specification (program) written in C or its "object oriented" version C++. Run the program through a "code profiler," which identifies performance bottlenecks. Analyze critical sections of code (the bottlenecks) to define custom instructions for faster execution. Translate the custom instructions into Tensilica Instruction Extension (TIE) language from which a custom microprocessor can be built in a foundry.

Design-time configurable microprocessors can be the individual blocks of a system-on-chip (SoC) design. Tensilica's average customer uses 6 microprocessor cores per design and some designs use more than 150. Moving up the level of abstraction, Tensilica's latest version of software can create superscalar and vector processors that are correct-by-construction, avoiding complex and costly verification. Using the high-level TIE language specification as input, the software profiles the program for bottlenecks and suggests a range of implementations, from minimum hardware to maximum performance. Ultimately, Tensilica will not need to describe its software in terms of microprocessor design at all; it will simply offer a range of "black-box" implementations to the user of its development software.

Design-time configurable microprocessors change the competitive landscape. These devices not only will invade markets of fixed-instruction-set microprocessors (such as PowerPCs, ARMs, and MIPS) but also can displace custom hardware (devices such as ASICS built and optimized to the utmost for a single function). Even a few years ago, I would have said this is impossible—that custom hardware is completely immune to assault by microprocessors. The microprocessor, after all, must use software to mimic the behavior of custom logic, so there seems to be no way it could be as efficient. But times change.

Combining flexibility and performance

The overthrow of custom hardware begins with new design methods. In the old days, when the engineer converted an English specification program, notoriously unclear, into a hardware description language, generally Verilog or VHDL, verification of correspondence among different levels of abstraction became an enormous task. For custom hardware, the verification task must be repeated for every application. For a microprocessor, verification occurs for only the original design and its cost is amortized across all of its applications.

In addition, Tensilica's configurable microprocessor bypasses verification of added custom instructions by tightly governing modification rules so that extensions are correct by construction. There's no translation of the executable specification into a hardware description language. Groups working on different blocks of a system-on-chip design all use the same standardized method (programming). Each function block of the SoC design can be a separate custom-configured microprocessor.

This advance combines the flexibility of a programmable microprocessor with the performance of a hard-wired custom circuit. Consider what custom hardware is; it is custom functions plus a state sequencer that controls the order and timing of execution step-by-step according to a clock. With standard manufacturer-designed microprocessors, there is a wide gulf between the fixed instruction set and custom-function capabilities. But with a configurable microprocessor, the base microprocessor is the state sequencer and the engineer can add an instruction that is the equivalent of a custom function, such as an encryption engine or a special communications filter. There's little difference between the performance capabilities of custom hardware and a custom-tailored microprocessor. They could be different words describing the same circuits. Configurable microprocessors blur the boundary between custom hardware and microprocessors.

In an era when cost-performance in desktop tethered systems gives way to cost performance *per watt* in mobile devices, custom chips might seem to have an edge in power efficiency. After all, a custom chip can forgo all power-using circuitry that is not needed for its one function. But a general-purpose microprocessor is assured of 10 or 100 times the unit volumes. Thus it can justify 10 to 100 times the design effort of a single-purpose custom chip. Microprocessor designers can spend effort on clock-gating, circuit-tuning, and power control that is way beyond what engineers can afford for custom hardware. Remember Tredennick's law and amend it: *Go for volume and you get performance* and *efficiency.* Recall my friend Joe and his four-stroke engine.

Stretch takes the next step

Tensilica's microprocessors are configured *at design time.* **Stretch, Inc.**, another Silicon Valley startup, takes the next step; its microprocessors are configurable *at run time.* Stretch is similar to Tensilica in several respects. It uses Tensilica's Xtensa microprocessor core as the central processor for its chips. Stretch's Chief Technical Officer, Albert Wang, was Chief Engineer at Tensilica. Though a startup itself, Tensilica is an investor in Stretch (lengthening the chain a little, **Altera** (ALTR) is an investor in Tensilica and I'm an investor in Altera, giving me a (small) stake in all three).

Unlike Tensilica, Stretch has no recipe for creating a custom microprocessor. Stretch's microprocessors are all the same. While Tensilica offers intellectual property that engineers license to turn into custom chips, Stretch builds and delivers generic chips. Implementation sounds about the same as Tensilica's process. Begin with a specification written in C or C++. Profile the application to identify bottlenecks. Stretch's software then creates custom instructions that speed the execution of these critical sections. Instead of creating custom hardware that is built through a foundry, however, Stretch's software

TELECOSM TECHNOLOGIES

Advanced Micro Devices	(AMD)
Agilent	(A)
Altera	(ALTR)
Analog Devices	(ADI)
Broadcom	(BRCM)
Cepheid	(CPHD)
Chartered Semiconductor	(CHRT)
Corvis	(CORV)
Equinix	(EQIX)
Essex	(KEYW)
EZchip	(LNOP)
Flextronics	(FLEX)
Intel	(INTC)
JDS Uniphase	(JDSU)
Legend Group Limited	(LGHLY.PK)
McDATA	(MCDTA)
Microvision	(MVIS)
National Semiconductor	(NSM)
Power-One	(PWER)
Qualcomm	(QCOM)
Samsung	(SSNLF/SSNHY)
Semiconductor Manufacturing International	(SMI)
Sonic Innovations	(SNCI)
Sprint	(FON)
Synaptics	(SYNA)
Taiwan Semiconductor	(TSM)
Terayon	(TERN)
Texas Instruments	(TXN)
VIA Technologies	(2388.TW)
Wind River Systems	(WIND)
Xilinx	(XLNX)
Zoran	(ZRAN)

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)

PARADIGM PLAY: SOFTENING HARDWARE, HARDENING SOFTWARE

SEPTEMBER 17: 19.88, 52-WEEK RANGE: 17.43 – 26.82, MARKET CAP: 7.39B

The company said sales would be flat sequentially and would not meet previous guidance of 2-4% growth. It was a smaller revision than main rival Xilinx, who previously said revenue would grow 2-4% sequentially but now believes it will shrink by 5-7%.

Altera and Xilinx stocks usually track each other, and it is difficult to exploit market inefficiencies in this arena. Although Xilinx's revenue is almost 60% higher than Altera's and it has almost $500 million more in cash and investment revenues, Altera is slightly more profitable with an operating margin of 29.9% compared to Xilinx's 27.4%. Thus Altera's price-to-sales and price-to-earnings multiples are usually higher than Xilinx's.

Although the companies have often leap-frogged each other over the years, neither gaining a clear advantage on technology, their strategies may be starting to diverge. As described by Nick Tredennick in this month's GTR, Altera is pursuing a high-volume general-purpose path, while Xilinx is tending to add more custom features to its chips, thus proliferating parts and diffusing talent, energy, and manufacturing costs across more products. Although both companies are well-positioned, if we had to pick one of the two, we'd probably go with the company speeding fastest down the learning curve, Altera.

On September 13, the board of directors increased the number of shares authorized for repurchase to almost 22 million—or about $440 million worth at today's price. At $19.57, the stock is off some 27% from its high earlier this year and now has a trailing PE of 32.9 and a calendar 2004 PE of about 26, assuming sales remain flat for the remainder of the year.

XILINX (XLNX)

PARADIGM PLAY: PIONEER OF PROGRAMMABLE LOGIC

SEPTEMBER 17: 28.54, 52-WEEK RANGE: 25.21 – 45.40, MARKET CAP: 9.91B

The company lowered expectations, citing inventory builds in Asia and domestic softness. Previous guidance of 2-4% sequential revenue growth gave way to an expected 5-7% revenue drop. Inventory days will approach 155, up significantly from the prior expectation of 135-140 days and even further from the company's target of 120 days, thus increasing the possibility of future product cuts or inventory write-downs. An apparent late-summer economic lull was felt across the semi space, with Intel, National, LSI, Broadcom, and competitor Altera also mildly paring expectations.

After peaking near 45 early in the year, Xilinx now trades below 28, a 38% fall off, with a trailing price-to-earnings multiple down to 28. Long term, Xilinx is still well positioned, and this softspot creates a buying opportunity. The company has a forward PE of 23, similar to others in the sector, but remains, we believe, positioned for stronger than average secular growth into new markets across the communications, industrial, and consumer electronics landscape.

In early September, Xilinx formed a new digital signal processing division, called Xtreme DSP, to tackle the growing wireless sector. Although Xilinx has already made inroads into DSP territory over the last few years, the company says there is a $2 billion opportunity to displace current ASIC and ASSP implementations and will now pursue the DSP market more formally. Headed by Omid Tahernia, a 20-year Motorola engineer and executive, the DSP division will also target the video, imaging, and aerospace markets, where high-speed real-time processing is required.

The company also named Mark Aaldering to head a new Embedded Processing Division, which will focus on Xilinx's Microblaze, Picoblaze, and PowerPC hard-core embedded microprocessors. Although Xilinx is likely to have some success with some customers, increasing hardware customization of its chips takes Xilinx away from the traditional and inherent advantages of general-purpose programmable logic devices.The *GTR's* Nick Tredennick thinks Altera's Nios soft-core approach can better deliver the promises of the high-volume PLD model.

QUALCOMM (QCOM)

PARADIGM PLAY: AIR KING—WORLD'S BEST TECHNOLOGY COMPANY

SEPTEMBER 17: 38.83, 52-WEEK RANGE: 20.50 – 41.17, MARKET CAP: 63.21B

Qualcomm has avoided the recent semiconductor blues and continues to trade at its highest level since early 2001. The company continued to push the strategic envelope with its early September $170-million purchase of Iridigm, a designer of special low-power color displays. Based on MEMS and thin film optics, Iridigm's iMoD displays could compete with or replace the power-hungry liquid crystal displays (LCD) used in most new phones, MP3 players, and other mobile devices. Iridigm says in addition to ultra low power characteristics, which could dramatically increase battery life, its iMoD displays are cheaper to manufacture than LCDs because they utilize a subset of the existing LCD fabrication infrastructure but require fewer process steps.

Iridigm's displays could add to Qualcomm's increasing dominance of most of the key components of mobile devices, including radios, baseband processors, applications processors, and software.

MEAD'S ANALOG REVOLUTION

NATIONAL SEMICONDUCTOR (NSM)
SYNAPTICS (SYNA)
SONIC INNOVATIONS (SNCI)
FOVEON
IMPINJ
AUDIENCE INC.
DIGITALPERSONA

COMPANIES TO WATCH

ATHEROS
ATI TECHNOLOGIES (ATYT)
BLUEARC
COX (COX)
ENDWAVE (ENWV)
LINEAR TECHNOLOGY (LLTC)
LUMERA (LMRA)
ISILON
MEMORYLOGIX
NOVELLUS (NVLS)
POWERWAVE (PWAV)
TECHNOLOGY
SEMITOOL (SMTL)
SIRF
SOMA NETWORKS
STRETCH INC.
SYNOPSYS (SNPS)
TENSILICA
XANOPTIX

Do We Love Lucent Now?

Now that Lucent Technologies (LU) is a leading supplier of CDMA and 3G infrastructure, some of our readers have asked us if are we ready to re-embrace this former favorite of our Telecosm list. Perhaps we can even celebrate the company as evidence for a return of the boom days of yore. After all, over the last two quarters it has claimed $664 million of operating income.

On the GTR list of telecosm companies from the opening gate, we touted Lucent in the late 90s for its apparent commitment to WDM and for its first-to-market broadband fiber (AllWave). After its sale of many of its paradigmatic technologies (including its fiber business) to raise desperately needed cash, we dropped Lucent in 2001. Lately, however, the company seems to be recovering steadily in its remaining wireline, services, and wireless businesses, managing to turn an operational-cash-flow loss of $1,900m/yr in June 2003 into a cash-flow gain of $440m/yr in June 2004. Wow. Since annual revenue has been flat during this period at about $8.7b, where's the cash coming from?

Not from services, which includes mostly low-margin product installation and development and which contributed just 16% to operating income in the June 2004 quarter (compared to 182% a year earlier when everything else produced losses). Nor from wireline, where revenues have decreased from $812m in the June 2003 quarter to $715m in the current quarter. With an operating margin of 9.5%, wireline contributes just 14.3% of operating income.

Which means the winner is … wireless, where quarterly revenue over the past year increased from $624m to $986m (or from 31% to 45% of total revenues), with operating margin increasing from minus 7.9% to plus 33.4%. Wireless now contributes a whopping 70% to Lucent's operating income.

So, now we love Lucent, right? Well, we still have to answer the question: Is Lucent primarily a paradigm investment or is it a bet on something else? We applaud the developing CDMA scenario, but like many dramas, the Lucent story has more than one act.

With the curtain barely opening on the second act, the news is already bad: Lucent's retiree benefits and pension plans are under-funded by about $7b. Based on financial and actuarial vagaries, projected discount rates and retirement schedules, these estimates are about as reliable as multi-decade weather forecasts or Keynesian economic astrology. Based on the benefits accrued as of today, the account is overfunded, even after cash payments to current retirees. So Lucent isn't paying anything in to the fund despite the portentous future projections.

So, maybe we, along with Lucent, should just pretend the future isn't coming. But we can't do that, because the future is *now* in the form of an accounting trick. Since the current account balance is positive, Lucent gets to add fictitious (noncash) pension credits to its operating results. Remove the bogus bucks, and the latest quarter's operating income drops from $349m to $72m, and March's sinks from $315m to a loss of $38m. Uh, oh, there goes our Telecosm celebration. Despite its claim to the contrary, Lucent is still losing revenues and struggling at operational breakeven. Unfortunately for Lucent, unless markets turn bullish very soon, the pension credits will begin to vaporize. That's because expected returns on post-retirement benefits are partly based on past-year returns, which still include the last bull market. The last bull will start to exit the equation shortly.

Forgoing applause for Lucent's second act and its anerobic financials, we now raise the curtain on Lucent's remaining legacy products. Wireless still contributes less than half of total revenues, with a third going to wireline, where steep declines in circuit-switching revenues continue. To fight these declines, Lucent is hoping for a revival in ATM sales, progress in VoIP softswitch revenues with the recent acquisition of Telica, increased sales in network-edge switches, and major contract wins in optical networking and switching where its LambdaXtreme and LambdaUnite compete with Ciena's CoreStream optical networking system and CoreDirector optoelectronic switch. Lucent recently beat out Ciena at Verizon where it will build the RBOC's next-generation long-haul network.

With the curtain still up on this third act, we ask if Lucent can beat the clock on legacy losses. Lucent may have answered the question for us: The company projects that revenues over next year or so should remain steady or increase slightly by about 5%, that gross margin will be flat to down from its current low 40s to the upper 30s, and that operating margins should remain flat to down by as much as 5% (meaning operating losses after subtracting the pension credits).

For the finale, we tackle valuations. In the short- and mid-term, liquidity is not an issue; working capital for the past year has held steady around $2.5b and reported book value is almost $10b (not counting pension costs), consisting mostly of tangibles and long-dated debt maturities. However, add in postretirement-benefit liabilities and book value falls to negative $3b. Similarly, the recent share price of $3.34 is a reasonable 16x estimated fiscal 2004 earnings, until you subtract the phony pension credits and discover that earnings disappear. Substituting the enterprise-value-to-sales ratio for PE, we get a lofty 3.5 compared to the average of 2.2 for Lucent's wireline competitors; for Lucent, an EV/S ratio of 2.2 is equivalent to a stock price of 67 cents.

Now, to answer our question: What does an investment in Lucent mean? It means you are betting that CDMA, VoIP switches, and semi-optical networking will win out over bogus accounts, legacy equipment, and languid bureaucracy. If you want to bet on CDMA, why not go to the best technology company in the world (Qualcomm) and invest without the risk of legacy liabilities and old-world thinking? Then again, if you want to invest in next-generation networks, Corvis is truly all-optical and comes without vagaries and bogus credits. And if you want to invest in VoIP, call your local cable company and try it out for yourself.

—Charles Burger

CONTINUED FROM PAGE 3

creates a personalization file on a memory that programs the configurable fabric of its generic chip.

Building a custom Tensilica chip requires design effort and expense for each application. Stretch's design effort and expense are amortized across the range of its applications. There's per-application effort in customizing the chip in the field, but this is perhaps an order of magnitude less than the expense to build custom instructions into a Tensilica chip, and it doesn't require expertise in hardware design. The tradeoff is that Tensilica builds execution support functions directly in custom hardware, while Stretch builds the functions in on-chip programmable logic. Programmable logic is necessarily slower than what could be provided for a Tensilica microprocessor. The consolation is that Stretch's configurable microprocessors offer performance that is way better than that afforded by standard (fixed-instruction-set) microprocessors.

I expect Stretch to be the first major commercial success for configurable systems

I expect Stretch to be the first major commercial success for configurable systems. Earlier efforts at configurable or reconfigurable systems required labor-intensive cooperation between engineers with detailed hardware knowledge and programmers with application knowledge. Stretch, by contrast, sells a straightforward, microprocessor-based programming model to programmers (no logic designers required).

Stretch's microprocessors are generic at manufacture and are customized in the field. The microprocessor's design and manufacturing costs are amortized across all applications, which reduces cost. Stretch's market will be applications that were outside the performance range for fixed-instruction microprocessors, but that don't have the volume to justify custom hardware design expense. It will also include applications with volumes to justify custom hardware expense, but that could not afford to forfeit the microprocessor's flexibility or that need quicker time to market than could be achieved with a custom hardware design.

The next step for Stretch? A real-time reconfigurable microprocessor. Personalization files, rather than being loaded when the system is initialized, could be loaded on demand at any time.

Altera and Xilinx: competing titans

Meanwhile, following in the footsteps of the startups and giving their innovations even more significance for investors and customers alike are the reigning giants of programmable logic devices (PLDs)—Altera and **Xilinx** (XLNX).

Altera has three major families, Max, Cyclone, and Stratix. Max chips consolidate 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.

Cyclone and Stratix chips, commonly called field-programmable gate arrays (FPGAs), have an associated SRAM personalization memory that is loaded each time the power comes on. The Cyclone family aims at low-cost, high-volume reprogrammable applications. The Stratix family aims at high-end, high-capacity applications.

The Xilinx families that correspond to Max, Cyclone, and Stratix are CoolRunner, Spartan, and Virtex, respectively.

Why do I mix discussion of these giant programmable logic vendors into a discussion of embryonic vendors of future microprocessors? Because, even though Altera and Xilinx today think of themselves as programmable logic vendors, they are becoming major vendors of microprocessors. To simplify the design process for programmable logic applications, they already offer on-chip hardware macros and extensive libraries of soft macros (commonly called intellectual property or "IP"). Among the most popular on-chip macros and IP functions are microprocessor cores.

In cores, Altera and Xilinx have diverging strategies. While both companies offer soft-core microprocessors, Altera offers no hard core versions of its popular Stratix or Cyclone FPGAs. By contrast, Xilinx's Virtex II Pro family offers chips with from one to four hard PowerPC cores.

A soft microprocessor core is a logic description that results in a microprocessor built in the chip's own fabric of programmable logic. A hard microprocessor core is custom hardware embedded with the programmable logic, but designed as a separate function. Compared to a soft core, a hard microprocessor core has better performance in less silicon area.

Configurability improves performance

Why would Altera abandon these advantages of hard cores? Offering chips with and without hard cores doubles the number of chips in the family. Parts proliferation increases per-chip cost because more unique parts mean lower per-part volumes and, therefore, higher per-part costs. So the hard cores should be on all chips or on none. But putting hard cores on all chips means some customers will pay for chip area they don't want. Also, for a given logic capacity, addition of a hard microprocessor core increases chip size. Because the cost of a chip increases exponentially with chip size, a chip that's a little larger costs a lot more. A 10 percent size increase, for example, might double a chip's cost.

Hard cores have additional costs. The custom microprocessor must be designed to match the semiconductor process for the programmable logic family. This means the microprocessor core offered on a programmable logic device will be a year or two behind the microprocessor's leading-edge implementation. For the most popular microprocessor cores,

such as ARM, MIPS, PowerPC, and Tensilica, there will be license fees and per-chip or even per-core royalties.

A hard microprocessor core is not configurable, but a soft core can be. Altera's Nios II soft-core microprocessor, for example, is configurable. It can implement custom hardware to support custom instructions. Nios II can have between 600 and 1,800 logic elements, comes in three compatible versions—economy, standard, and fast, and commands a library of 60 peripherals—all for a logic cost of from 1 to 15 percent of a low-end Cyclone FPGA. Based on chip cost and on its use of the chip's resources, the cost of a soft-core micro is from thirty-five cents to less than three dollars, which is competitive with the costs of low-end, stand-alone microprocessors. Use of the Nios microprocessor is royalty-free, but requires the one-time purchase of a $495 development kit. (As development kits go, this is pretty close to free.) As a configurable device, the Nios microprocessor offers up to 256 custom instructions that can be added to the base set.

Configurability not only improves performance; it also enables novel applications.

For example, **Motorola**'s (MOT) Canopy fixed wireless access system (for "last mile" Internet access) uses a single configurable hardware design for both a base station and a client. Instead of building an expensive, low-volume base station and a low-cost, high-volume client, Motorola built a single, configurable high-volume node. It is configured in the field as either a base station or as a client. One design effort, one bill of materials, one part number in the supply chain—all leading to lower cost.

The Studio MovieBox Deluxe from **Pinnacle** Systems (PCLE) is another example. The MovieBox imports almost any digital or analog video source, whether VCR, DVD, web cam, TV or camcorder, through a USB connection into the PC. As a USB, consumer-PC peripheral, it needed both low cost and low power. To get low power, the designers threw out the microprocessor and, for each function, designed custom logic to implement in an Altera Cyclone FPGA. Translation functions are programmable logic personalization files which are a part of the PC's MovieBox device driver. In the MovieBox itself there's no microprocessor and there's no memory. How's that for a novel way to minimize component cost and to conserve power?

Strategic positioning

The programmable logic market is about $3 billion, but it is growing into the $30-billion ASIC market and into the $40-billion microprocessor market. The programmable logic companies know how to sell to the ASIC market because ASIC application development is almost identical to programmable logic application development. The customers are logic designers. Moore's-law progress makes programmable logic devices suitable for ever-larger segments of the ASIC market. Today's mid-range devices from Altera and from Xilinx have 200 times the capacity (in logic elements) and they run 40 times faster than the devices of ten years ago.

The market for traditional microprocessors and for DSPs differs from the ASIC market: the crucial customers are programmers, not logic designers. This market is also the market target for ARC, MIPS, Stretch, and, with the latest releases of its software, Tensilica. These companies already sell either directly to programmers or they are moving in that direction. To be competitive, Altera and Xilinx will have to reorient their businesses to sell to programmers. That means reorienting the development software and the marketing message from logic designers to programmers.

For the last twenty years, Altera and Xilinx have grown rapidly with blooming demand for their products. In listening to and in serving their customers, they have been going "up market," emphasizing high-end, high-margin chips. For Altera, that's been Flex, Apex, and Stratix; for Xilinx, it's been the X2000, X3000, X4000, and Virtex. With the industry's transition in emphasis from tethered to mobile and from industrial systems to consumer systems, a high-end strategy will paint these manufacturers into a corner of diminishing returns. (A $30-billion ASIC market above them gives them some but not unlimited time.)

Whether or not by a deliberate decision on Altera's part, the success of its low-end Cyclone family and its concentration on the soft-core Nios now give it the best strategic position to move down market. Build for volume and performance follows. Cyclone family chips will grow in the market in three areas: low-cost, design novelty, and more logic capacity.

At the low end, Cyclone will displace "mature" older-generation components. Many of these components are immune to Moore's law advances because the chip size is determined by the chip's connections to the outside world. For an increasing number of chips, the connector pad ring, rather than the chip's logic circuits, sets its size. Once that happens, the chip stops shrinking, which means it stops getting cheaper (cost is fixed by chip size). Once the pad ring sets the chip's size, whether it's custom hardware or programmable logic, the cost is the same. That gives programmable logic the advantage because a programmable generic component will have higher volumes and thus lower costs and more flexibility than the application-specific chip.

See-saw battle

With growing logic capacity, the Cyclone family will tend to push Stratix into an up-market corner. Even so, Altera cannot abandon Stratix because the high-end components have handsome margins, needy customers, and a critical role in maintaining a leading-company reputation. Further, Stratix chips serve as the "concept cars" for Cyclone in the same way that experimentation in high-performance system design by the workstation companies benefited the PC.

So who will win the long-running battle between Altera and Xilinx? As a former Altera employee, I'm biased in favor of Altera, but I hope neither wins. Competing titans move

the industry forward faster and better than virtual monopoly does. It's easy to imagine how the situation in the PC industry might have been improved had there been parity competition for **Intel** (INTC) in microprocessors and for **Microsoft** (MSFT) in applications and operating systems.

Altera and Xilinx have been in a see-saw battle. Lately, Altera has refreshed its entire product line, improved public and press relations, and is disclosing its roadmap. It is well positioned at the low-end with Max, with Cyclone, and with the Nios II soft-core microprocessor. Nonetheless, Xilinx still has a substantial advantage in revenue and a huge advantage in mindshare among engineers. Because Altera does not have a CTO, Xilinx's CTO speaks for the entire programmable logic industry. Xilinx's mindshare in the engineering community comes from its superb university program, from its conference sponsorship and participation, from its strategic investments, and from its public relations.

There's still a large gulf between the performance and capabilities of an ASIC and its competitors (microprocessors and PLDs). Products called structured ASICs, from companies such as **Chip Express, eASIC, Faraday Technology, Lightspeed, LSI Logic** (LSI), and **NEC** (NIPNY), hope to fill this gap. But a large gap doesn't mean a large need. Most applications are at the low-end and a few (that must have ASICs) are at the high-end, leaving slim pickings in the unserved middle. Further, structured ASICs attack the cost of producing a custom chip, but design costs, which may be ten times as high, are largely unchanged.

I used to think that reconfigurable systems from Altera and Xilinx would emerge to displace microprocessors in power-sensitive and performance-oriented systems. In these systems, custom hardware would be "demand paged" into a chip of reconfigurable resources. But that's not how we'll get to real-time reconfigurable systems.

Microprocessors, as we are already seeing with ARC, Altera's Nios, MIPS, Stretch, and Tensilica, will first become configurable and compete with fixed micros in performance and with ASICs in flexibility and cost. Then will come the next step, perhaps from my favorite unfunded startup **Ascenium**.

Today's configurable microprocessors run a code generator that emits standard instructions for most program execution, using custom instructions and associated custom hardware only to ease bottlenecks. Ascenium's compiler substitutes a circuit generator for the code generator. An Ascenium microprocessor reconfigures its resources with each custom instruction to execute the equivalent of hundreds or thousands of standard instructions.

Change is in the air

The fixed-instruction-set microprocessor's success has built a $40-billion applications market in thirty years. Similarly, ASICs have built a $30-billion applications market. The markets grew separately because the microprocessor did not have the performance to challenge custom hardware applications and ASICs couldn't meet the low development cost of microprocessor-based implementations. The time for both is ending. In the microprocessor market, configurable microprocessors will displace DSPs and high-end fixed-instruction-set microprocessors. In the ASIC market, configurable microprocessors will rapidly overtake low-end applications. It will look as if the microprocessor market is growing at the expense of the ASIC market.

Stretch is well-positioned in the near term; Altera and Xilinx are well-positioned for the long term as they make the transition from being chip companies to being microprocessor companies. Makers of fixed-instruction-set microprocessors, of DSPs, of ASICs, and of structured ASICs will lose ground slowly because the total market will grow even as the configurable microprocessors invade. Since all of the providers of configurable microprocessors are fabless, foundries will gain, including **Taiwan Semiconductor** (TSM), **Semiconductor Manufacturing International** (SMI), **United Micro Electronics** (UMC), and **Chartered** (CHRT). Integrated device manufacturers such as Intel, and TI will tend to lose share unless they learn how also to perform as foundries.

The change portended by Tensilica's blow-away success in the benchmark tests will be slow, but the direction is inexorable. Investors should get on the right side of it. At the moment the best vehicles are Altera and Xilinx. But watch this space for IPOs.

– Nick Tredennick with George Gilder and Brion Shimamoto
September 16, 2004

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 ©2004, by Gilder Publishing, LLC

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

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