

StarCore Launches First Architecture

Lucent and Motorola Disclose New VLIW-Based Approach



by Ole Wolf and Jeff Bier,
Berkeley Design Technology, Inc.

Although Lucent and Motorola's joint DSP development center in Atlanta hasn't yet opened its doors, its first design has already been announced. With the new StarCore 400 architecture, revealed at Microprocessor Forum earlier this month, the design team aims high, promising stellar performance in nearly the entire DSP application space.

The StarCore 400 targets an unusually wide array of applications, ranging from low-power and cost-sensitive applications such as cellular handsets to high-performance infrastructure applications such as modem banks and cellular base stations. To enable support for such a vast range of applications, StarCore 400 defines a scalable architecture that can be adjusted to suit the needs of specific applications. Different implementations of the StarCore 400 architecture will sport varying constellations of execution units, buses, and other resources tailored for the application at hand.

Gather Ye Starlets

Lucent and Motorola announced their "collaborative R&D partnership" in June (see MPR, 6/22/98, p. 10). The focus of the partnership is the development of next-generation DSP cores. Lucent and Motorola will separately develop products based on the new cores, though they have held out the possibility of second-sourcing each others' products. The partners are also cross-licensing three existing cores: Motorola's

DSP56800 and M•Core microcontroller as well as Lucent's DSP16000. By working together, Lucent and Motorola hope to challenge Texas Instruments' durable dominance in DSPs. Beyond combining their resources to speed the development of new architectures, cores, and tools, the partners hope to create enough momentum behind their new architectures to attract the increasingly critical support of independent tool, board, and application-software developers.

Lucent and Motorola are each significant players in the DSP market in their own right, though their market shares are dwarfed by TI's, as Figure 1 shows. The bulk of Motorola's DSP sales (estimated at 70%) are to other Motorola divisions. In the merchant market, Motorola has lost traction, and its overall market share, including internal sales, has declined from 17% in 1993 to 12% in 1997. Motorola has been through numerous reorganizations in recent years, leaving some customers—and no doubt some employees—confused about its direction. The StarCore partnership may be the catalyst needed to give Motorola a crisper focus and clearer direction for its DSP products.

Lucent, in contrast, has been pursuing a fairly clear and consistent strategy—one of specialization. Lucent has focused its DSP products exclusively on telecommunications applications, and it has focused its sales efforts on the very largest top-tier equipment manufacturers. This strategy has worked well for Lucent, gaining it the number-two market share position despite a relatively narrow product line and limited third-party support.

But as the market for DSP processors broadens, Lucent may need to address more customers and perhaps more applications to maintain and grow its market share. A key ingredient in achieving this expanded appeal will be a broad base of high-quality tools, application software, and other third-party infrastructure—an area where TI enjoys unchallenged leadership today. By adopting architectures in common with Motorola, Lucent stands a much better chance of garnering the support necessary to serve a larger customer base.

Significantly, the partnership may at long last break the one-to-one mapping between architectures and vendors. At present, choosing a DSP architecture means choosing a chip vendor, since each supplier has its own proprietary architectures. Thus today, once users have chosen an architecture, they're limited to products from a single vendor, based on that architecture. This limits the selection of on-chip memory and peripheral configurations, packaging options, and so forth. If StarCore 400 succeeds, customers may instead be able to buy compatible devices from Lucent and Motorola.

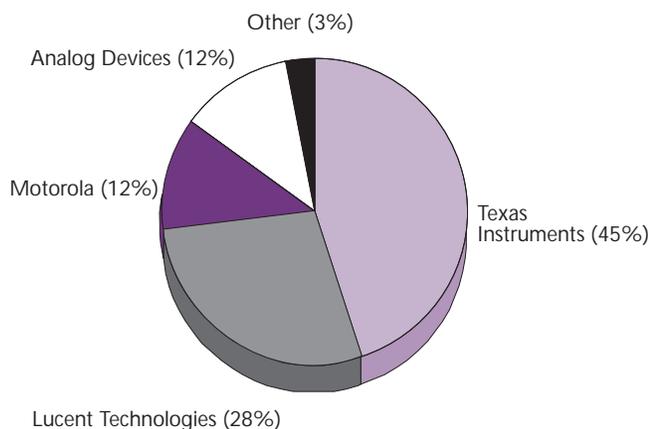


Figure 1. DSP merchant market share figures for 1997 show that Motorola and Lucent together nearly match TI. Motorola's share includes sales to other Motorola divisions, but Lucent's market share data excludes internal sales (believed to total less than 15% of Lucent's sales). These figures exclude hybrid CPU/DSPs. (Source: Forward Concepts)

This compatibility highlights a major risk factor for the success of the partnership, however. The partners will ultimately compete against each other, offering comparable products to the same customers. This competition, similar to the relationship between former PowerPC partners IBM and Motorola, may make it difficult for Lucent and Motorola to collaborate effectively over the long term.

When the StarCore partners announced their venture, they projected that their joint design center would open in the third quarter. The center is now expected to open in November, slightly behind schedule. Nevertheless, it is clear that the partners have been busy for quite some time on their first architecture.

Little Light Escapes

So far, Motorola and Lucent have disclosed few details about the StarCore 400 architecture and no information at all about specific products to be based on the new architecture. In his Microprocessor Forum presentation, Zvika Rozenstein, Motorola's chief architect for StarCore, described a VLIW-like architecture with an emphasis on scalability.

The StarCore team uses the term VLES, or variable-length execution set, to describe the style of instructions and the manner in which they are executed in the new architecture. Like TI's 'C6xxx devices, introduced in 1997 (see [MPR 2/17/97, p. 14](#)), StarCore 400 implementations will issue and execute a varying number of simple instructions per cycle. Instructions will be scheduled at compile time (or manually by assembly-language programmers) rather than at run time, with each instruction mapped dynamically to a specific execution unit at run time.

The StarCore 400 architecture uses 16-bit instructions, in contrast to the 32-bit instructions of the latest TI devices. This should give StarCore an advantage in code density, which is a weak point for TI's 'C6xxx. It isn't possible, however, to encode the full range of the operation and operand combinations supported by a powerful architecture using a 16-bit instruction.

For this reason, the StarCore architecture allows a varying number of optional "prefix" instructions to be included in each group of simultaneously issued instructions. The prefix instructions extend the function of the basic instructions—for example, allowing access to a larger number of registers or enabling predicated execution. Generally, the prefix instructions affect the entire group of parallel instructions that follow. This approach is illustrated in Figure 2.

Code Size Dwarf?

This variable-length instruction approach has worked well for some recent DSPs, such as Lucent's DSP16000 family, allowing them to combine the power of a 32-bit instruction word where it is needed (primarily in performance-critical inner loops) with the code density of a 16-bit instruction word. The StarCore team claims that StarCore 400 implementations will have better code densities than conventional

Single 16-bit instruction



Multiple 16-bit instructions



Multiple 16-bit instructions with 16-bit prefix instructions



Figure 2. StarCore 400 variable-length execution sets. In each instruction cycle, a variable number of 16-bit instructions are executed. Instructions can be extended to an effective instruction-word width of more than 16 bits, using "prefix" instructions.

DSP processors, comparable to those of the M•Core and ARM7 embedded controllers, as Figure 3 shows.

On the TI 'C6xxx, code size is increased because many operations have multicycle latencies, requiring software pipelining and loop unrolling to obtain peak performance. The partners have said very little about the pipeline of the StarCore 440—the first StarCore 400-based core—except to assert that it will be "short," and that operations such as multiply-accumulate will have single-cycle latencies. On the 'C6xxx, however, delayed loads and branches cause the greatest challenge for code density, and the StarCore team has said nothing about the latencies of these operations on their new architecture.

Based on the sketchy information disclosed, it isn't possible to confirm the code-density claims. That will have to wait until sometime in the first half of 1999, when the alliance promises to make StarCore 440 documentation and tools broadly available. According to the partners, selected customers are already using early versions of the tools.

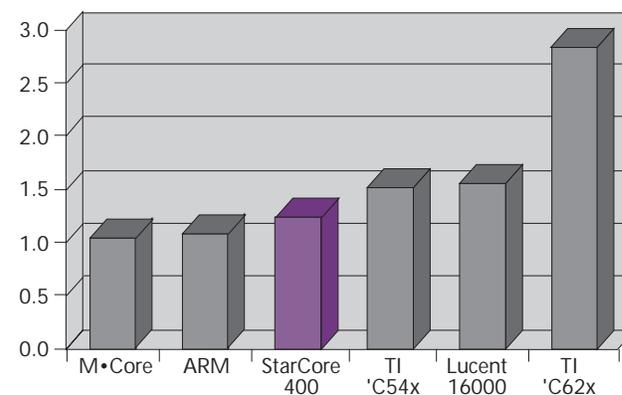


Figure 3. StarCore's benchmark results showing relative memory usage of the StarCore 400 architecture on compiled C and C++ code. According to Lucent and Motorola, these results are based on a subset of DSP, control, and encryption benchmarks drawn from Motorola's own "PowerStone" benchmarks.

For More Information

According to Motorola and Lucent, the first devices using the StarCore 440 core will begin sampling to customers in late 1999 or early 2000. Neither company has yet announced actual products based on the core. The StarCore Web site is <http://starcore-dsp.com>.

A Universal Architecture?

Lucent and Motorola are emphasizing the scalability of the StarCore 400 architecture, claiming that various manifestations of the architecture will be able to address an extremely broad range of applications, from low-power wireless communications devices to performance-hungry infrastructure equipment.

According to Lucent and Motorola, the StarCore 400 architecture is scalable in several dimensions to meet the needs of these widely varied applications. As examples of this configurability, the partners cite clock speed, data width, number and type of execution units, memory bandwidth, number of registers, and the types and maximum number of instructions that can be executed per instruction cycle. For example, for memory-bandwidth-intensive applications such as video processing, the width and number of the on-chip buses and the number of associated address-generation units and address registers might be expanded. Similarly, if an application lends itself to a high degree of parallelism, a wider instruction-issue bandwidth and more execution units can be employed.

In addition, the designers state that specialized execution units and registers can easily be incorporated to meet the needs of specific types of applications. Figure 4 reproduces the designers' conception of the key scalable elements of the architecture. Since StarCore 400 makes provisions for

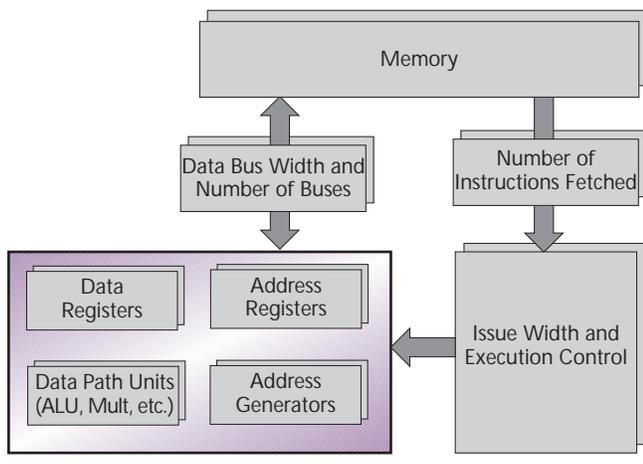


Figure 4. According to the StarCore partners, scalable features of the StarCore 400 architecture include data bandwidth, instruction issue width, function units, address generators, and registers.

changing programmer-visible state (such as the number of registers) from one implementation to another, some would argue that it isn't an architecture in the usual sense of the word. Lucent and Motorola have suggested, however, that different implementations of the StarCore 400 architecture will be upwardly object-code compatible, easing users' transitions from one implementation to the next. Nevertheless, reoptimization of existing code will be required to take full advantage of more powerful implementations.

It is not clear whether the StarCore 400 is intended only for fixed-point applications, or whether floating-point data types will also be supported. In recent years, both Motorola and Lucent have dropped their floating-point DSPs to focus on higher-volume fixed-point devices. There appears to be nothing in the StarCore 400 architecture, however, to prevent core designers from building a floating-point version, such as TI has done with its new 'C6701 (see MPR 9/14/98, p. 18).

According to the StarCore team, tools will be available in the first half of 1999. The tool chain will orbit around an integrated development environment and will include a C and C++ compiler, assembler, simulator, and source-level debugger and profiler. The compiler is responsible for scheduling instructions to take maximum advantage of inherent parallelism, which is not a simple task. In addition, an assembly-language optimizer—a new category of tool first offered by TI with the 'C6201—will be provided to assist assembly-language programmers in producing tight code. According to the designers, StarCore 400 on-chip debugging support will include nonintrusive host-target data transfers and event-tracking capabilities.

The scalability of the architecture, if it is realized, will present significant challenges for tools developers, who must create tools flexible enough to efficiently support a range of StarCore 400 implementations. Nonetheless, the StarCore team makes aggressive claims about compiler performance. According to StarCore, its architecture will be more amenable to compiler code optimization, due to its short pipeline and a relatively homogeneous array of execution units.

The team states that its compilers, which incorporate "new technology" from unnamed third parties, have evolved with the StarCore 400 architecture and will effectively detect parallelism and make efficient use of the processor's resources. In the DSP universe, truly efficient compilers are rarer than Halley's comet, which suggests that skepticism is in order until StarCore delivers. If it does deliver, however, the StarCore team will have a significant competitive advantage. Even if the code-generation tools fall short, the shorter pipeline and more homogeneous architecture should simplify the work of assembly-language programmers tasked with creating highly optimized application code.

How I Wonder What You Are

The joint venture partners will reportedly begin sampling their first devices based on the new architecture in late 1999

or early 2000. Unfortunately, the StarCore team has been nebulous regarding the details of the 440, perhaps in an effort to keep the competition off balance. As of this writing, the designers have disclosed that the core uses 16-bit fixed-point data paths with 40-bit accumulators and executes up to six instructions per cycle, including up to four multiply-accumulate operations. In addition, the core has a 4G, byte-addressable, unified address space (i.e., instructions and data occupy the same address space) and a data-memory bandwidth of eight 16-bit words per cycle.

Although the StarCore team seems intent on keeping observers in the dark about most details of the StarCore 440 implementation, the team has made very specific performance claims. The core will be implemented in Motorola's HIP6 process, which Motorola calls a 0.13-micron process, though it is comparable to other vendors' 0.18-micron processes. In the HIP6 process, the initial StarCore 440 device has a projected clock speed of 300 MHz, yielding a maximum of 1.2 billion multiply-accumulate operations per second. With a 1.5-volt supply, power consumption of the core plus an unspecified amount of on-chip SRAM is projected to be under 180 mW at 300 MHz.

In a sense, StarCore's adoption of a VLIW-based architecture validates TI's choice of a VLIW approach for its 'C62xx family. At the same time, StarCore 440 presents a strong challenge to TI. If the new core delivers on its promises, the devices using it will be noticeably faster than current 'C62xx devices. More important, however, they will deliver this speed with much lower memory usage and power consumption, and with reduced code-generation complexity.

With the StarCore 440, the StarCore team thus promises speeds much higher than those of today's DSPs—but without penalties in the form of high memory usage and power consumption. If the team succeeds, StarCore 440 will be a potent competitor to most commercially available DSPs. It is inevitable, however, that in the one- to two-year period it will take to roll out the first StarCore products, competitors will make major advances as well.

Deep Impact Possible

The StarCore partners face a number of serious challenges in bringing StarCore-based devices successfully to market. First, making the partnership itself work effectively will not be easy. While high-profile joint technology-development efforts are common among large companies, the number of such alliances that have yielded successful products is remarkably small. The StarCore partners will need discipline and a measure of luck to combine the best of what both companies have to offer, rather than the worst, and to

continue cooperating even as they compete to sell StarCore-based devices.

Architecturally, StarCore 400 represents a significant course change for both Motorola and Lucent, both of which have heretofore relied exclusively on fairly traditional DSP architectures. In addition, the emphasis on scalability at the architectural level is a first for DSPs. The partners' goal of serving an extremely wide range of applications with a single architecture is quite ambitious, and reaching it will be challenging. If the partners succeed, however, the rewards should be commensurate with the effort.

Clusters Form

At the Microprocessor Forum, Analog Devices, the fourth major DSP competitor, also unveiled a new VLIW-based architecture. Dubbed TigerSHARC, the new architecture represents the third generation of ADI's SHARC floating-

point devices. The ADI and StarCore announcements complete a sea change in DSP architectures. Whereas just two or three years ago all of the major DSP vendors seemed firmly entrenched in conventional single-issue DSP designs, now all four major players have committed to VLIW-like approaches to varying degrees.

But TI and ADI appear to consider their newest architectures best suited for high-end applications, whereas the StarCore partners are pushing the VLIW approach for more cost-sensitive and power-conscious designs. It will be interesting to see whether other DSP developers—for example, DSP core suppliers and Japanese semiconductor vendors—soon join the growing VLIW bandwagon.

As with any new architecture—and especially one targeting a broad range of applications—tools will be critical to the success of StarCore. Tools for DSP-based applications have traditionally been a weak point. As architectures become more complex and powerful, and applications larger and more demanding, the importance of high-quality tools increases significantly. In addition to the baseline software tools being developed by StarCore, the partners must convince a significant number of third-party tool, board, and application-software vendors to support the new architecture.

If the StarCore partners can keep their relationship productive, deliver the promised performance on time, and garner the necessary third-party infrastructure support, they will have dramatically improved their footing for competing successfully with TI in the next decade. ■

Authors Ole Wolf and Jeff Bier are with Berkeley Design Technology, Inc., (www.bdti.com) a DSP technology analysis and software development firm. Wolf and Bier are co-authors of Buyer's Guide to DSP Processors, the 1999 edition of which will be available from MDR shortly.



Zvika Rozenshein, chief architect for StarCore, describes the VLES style of instruction-set design.

MICHAEL MUSTACCHI