

SiByte Reveals 64-Bit Core for NPUs

Independent MIPS64 Design Combines Low Power, High Performance

By Tom R. Halfhill {6/26/00-04}

On June 12, unfazed by the burgeoning number of network processors (NPUs), SiByte disclosed the first details of its new SB-1 microprocessor core at **Embedded Processor Forum**. If the Silicon Valley startup can deliver what it promises—a 1GHz core that surpasses

2,000 Dhrystone mips while consuming only 2.5W—the SB-1 will push MIPS-based NPUs to new heights of power efficiency and performance.

Great things are expected of SiByte because of its pedigree. The president and CEO is Dan Dobberpuhl, who was responsible for designing the Alpha, StrongARM, and MicroVAX processors at Digital Semiconductor in the 1990s. Since cofounding SiByte in July 1998, Dobberpuhl has assembled a team of more than 90 engineers, including many refugees from Digital. Last year, SiByte licensed the MIPS64 instruction-set architecture (ISA) from MIPS Technologies and began working on an independent core design (see [MPR 12/27/99-en](#), “SiByte Licenses MIPS for Network Processor”). SiByte plans to use the new SB-1 core in a family of integrated NPUs for the networking and communications markets.

Of course, that plan puts SiByte on a collision course with numerous other companies pursuing the same goal. A brief list of the higher-profile competitors includes IBM, with its Rainier chip (see [MPR 10/6/99-en](#), “IBM, C-Port Network Processors Challenge Intel”); Intel, whose IXP1200 was also designed by former Digital engineers (see [MPR 9/13/99-01](#), “Intel Network Processor Targets Routers”); Lexra, which also announced its NetVortex core at Embedded Processor Forum; Motorola, which recently acquired C-Port (see [MPR 3/6/00-en](#), “Motorola Buys C-Port: Smart Move”); and Sitera, recently acquired by Vitesse (see [MPR 5/29/00-02](#), “Sitera Samples Its First NPU”). All those companies—and more—

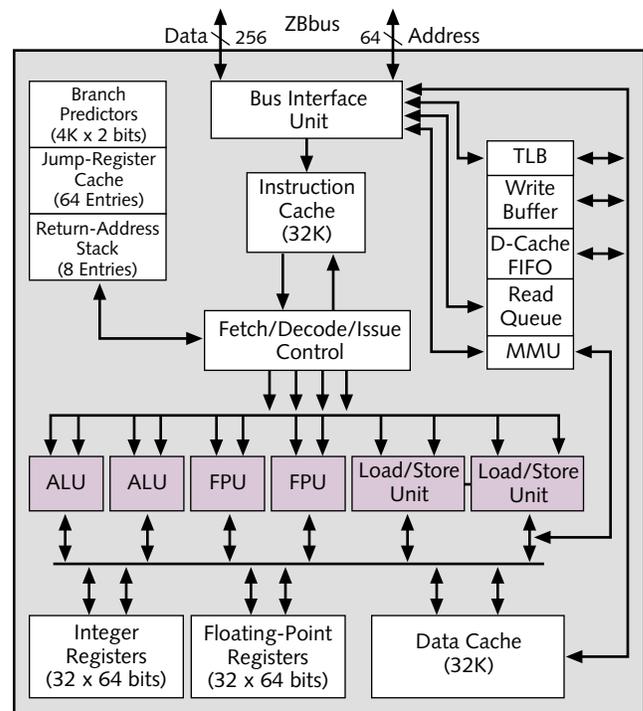


Figure 1. SiByte's SB-1 core is a four-issue superscalar design with six function units, shown here in purple. Note the extensive resources for branch prediction (upper left) and the wide bus interface (top), which supports chip multiprocessing.

are introducing programmable NPUs designed to replace ASICs or general-purpose CPUs in the packet-processing paths of routers, switches, and other networking equipment.

Fortunately for SiByte, the SB-1 is a highly competitive design. It will probably exceed the performance of MIPS Technologies' own 20Kc, another MIPS64-based core announced at Embedded Processor Forum. The SB-1 has twice as much execution bandwidth and is expected to run at clock frequencies at least 33% higher than the 20Kc without consuming more power. And, like Lexra's NetVortex, the SB-1 has special features for chip multiprocessing (CMP).

SiByte plans to deliver integrated chips based on the SB-1 core in 2001. That's roughly in line with the MIPS and Lexra plans for rolling out the 20Kc and NetVortex cores. However, SiByte's first NPUs will follow similar chips from IBM, Intel, C-Port, and Sitera by 12 to 18 months.

Vector Instructions Boost Throughput

As Figure 1 shows, the SB-1 is a four-way superscalar design with a well-balanced array of function units. It can issue two integer or floating-point instructions in parallel with two load/store instructions on every clock cycle. The FPUs support single- and double-precision IEEE-754 floating-point operations with a typical latency of four cycles and single-cycle throughput. In terms of both features and performance, it's the most powerful core ever built on the MIPS64 architecture.

Under ideal conditions, a single SB-1 core running at its maximum target clock speed of 1GHz can execute 8 bops (with 16-bit integers) or 8 gflops (with 32-bit floats). Vector instructions in the MIPS64 architecture make those peak execution rates possible. To reach 8 gflops, each FPU executes a vector multiply-add (MADD) instruction that operates on a pair of 32-bit values in a 64-bit floating-point register. These "paired-single" instructions are essentially the same as the single-instruction multiple-data (SIMD) instructions in other CPU architectures. They first appeared in the MIPS

architecture with the MIPS-V instruction set and were joined by additional paired-single instructions in the MIPS-3D extension (see *MPR 8/23/99-en*, "MIPS Adds a New Dimension to MIPS64").

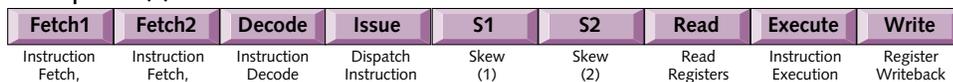
Such powerful floating-point capabilities might seem out of place in a core designed for NPUs, because packet processing relies heavily on bit-manipulation and other integer-type operations. But routing isn't the only intended application for the SB-1. SiByte also hopes to sell chips for devices that need DSP capabilities, such as broadband modems, access concentrators, and wireless base stations. Also, network devices that handle secure traffic on virtual private networks and commercial Web sites must encrypt and decrypt the data packets, and that's an FPU-intensive task.

SiByte plans to introduce SB-1 chips at clock frequencies centered around 800MHz and quickly ramp to 1GHz in a 0.15-micron CMOS process from TSMC, SiByte's foundry partner. At 1GHz, according to SiByte's simulations, the core will deliver at least 2,000 Dhrystone 2.1 mips and consume about 2.5W at 1.2V. The estimated die size is 25mm²—including the 32K instruction and data caches, which are four-way set-associative.

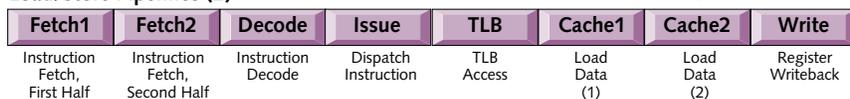
One way the SB-1 attains those high clock frequencies is with deeply pipelined function units. As Figure 2 shows, the ALUs have 9-stage pipelines, the load/store units have 8-stage pipelines, and the FPUs have 12-stage pipelines. In comparison, the MIPS 20Kc has two 7-stage integer pipelines and a single 9- or 10-stage floating-point pipeline (the 9-stage path is optimized for MADDs).

Two special "skew" stages in the SB-1's ALU and FPU pipelines allow the processor to execute a load instruction immediately followed by a dependent integer or floating-point instruction without creating any pipeline bubbles. The skew stages give the load instruction enough time to fetch a pair of operands from the data cache and forward them to the dependent integer or floating-point instruction. Without the skew stages, the dependent instruction would

ALU Pipelines (2)



Load/Store Pipelines (2)



Floating-Point Pipelines (2)

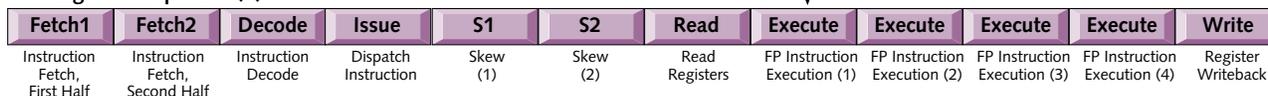


Figure 2. The SB-1 is deeply pipelined, and the ALU and FPU pipes have special "skew" stages that allow a dependent instruction to follow a load without stalling. In this diagram, the arrows show how load instructions can forward data to the other pipelines in synchronization with their execute stages.

stall for a couple of cycles while waiting for data to arrive, creating bubbles or empty stages in the pipeline. The trade-off for this feature is a longer execution pipeline, which levies a greater branch-misprediction penalty. A processor that allows out-of-order execution could avoid this problem by issuing a nondependent instruction after the load, but the SB-1 is strictly an in-order machine.

Deep pipelines are especially vulnerable to bubbles caused by taken branches, so the SB-1 has unusually sophisticated branch prediction for an embedded processor (or even for a PC processor). Dynamic predictions are based on 2-bit Gshare history tables (strongly taken, weakly taken, weakly not taken, strongly not taken). A direct-mapped cache stores 4K of those entries. A 64-entry jump-register cache stores branch-target addresses, and an 8-entry wraparound stack stores return addresses. Judging from the effectiveness of similar mechanisms in other types of microprocessors, the SB-1 should be much better than 90% accurate at predicting the directions of branches. A factor in its favor is that packet-processing software generally doesn't branch as unpredictably as user-directed PC software.



Robert Stepanian, director of architecture at SiByte, describes the four-issue superscalar core of the SB-1.

Designed for System-On-a-Chip CMP

Some other salient features of the SB-1 are a translation-lookaside buffer (TLB) that stores 64 pairs of entries; an 8-entry read queue that allows the nonblocking data cache to handle eight outstanding cache misses; a 16-entry FIFO stack that holds pending stores to the data cache; and a 4-entry write buffer.

To support multiple cores on a chip, the SB-1 has a fully coherent split-transaction bus interface referred to as the ZBbus. This bus has a 256-bit data path with 64-bit addressing, and it runs at half the processor's core frequency. At 500MHz, its peak data bandwidth is 16GB/s. As Figure 3 shows, the ZBbus allows several SB-1 cores to share the same secondary cache, memory controller, and I/O bridges to on-chip peripherals and peripheral buses.

The ZBbus can accommodate almost any number of SB-1 cores, but the manufacturing process imposes a practical limit. In a typical 0.15-micron process, a complement of two to eight cores is a reasonable scale of integration. To maintain memory coherency among multiple cores, the SB-1's bus-interface unit stores a duplicate set of data-cache tags and allows other cores on the ZBbus to snoop the tags. MIPS-standard load-link and conditional-store instructions allow multiple cores to synchronize their atomic read-modify-write operations.

SiByte engineers optimized the SB-1 for high performance and low power consumption by using synthesis tools

sparingly and only for noncritical blocks. All critical paths benefit from full-custom circuit layouts derived from hand-drawn schematics. SiByte says it cut power consumption 36% by powering the core at 1.2V instead of 1.5V, another 30% by using clock-gated circuits extensively throughout the design, another 10% by using flip-flops instead of latches, and another 5–10% by selectively using dynamic logic (the core is mostly static).

Careful attention to circuit design was the main reason that StrongARM, Dobberpuhl's earlier project at Digital, achieved an exceptional combination of high performance and low power consumption. At SiByte, however, his engineers are working under a handicap—they cannot design as closely to the IC process as the StrongARM engineers at Digital could. Digital owned a fab in Hudson, Massachusetts, where it manufactured StrongARM chips in an optimized 0.35-micron process. (In fact, StrongARM chips are still manufactured in that same process by the fab's current owner, Intel.) SiByte, a fabless company, must rely on independent foundries for manufacturing.

As a result, the SB-1 adheres to TSMC's generic 0.15-micron design rules. Dobberpuhl acknowledges that a fabless business model has slightly compromised the design of the SB-1, but he says the core achieves an exceptional power/performance ratio nevertheless.

He's not making an idle boast. As Table 1 shows, the SB-1 competes strongly against some other recently announced cores for high-performance embedded applications. If NPUs based on the SB-1 perform as well as SiByte's simulations indicate, the chips will be formidable competitors in

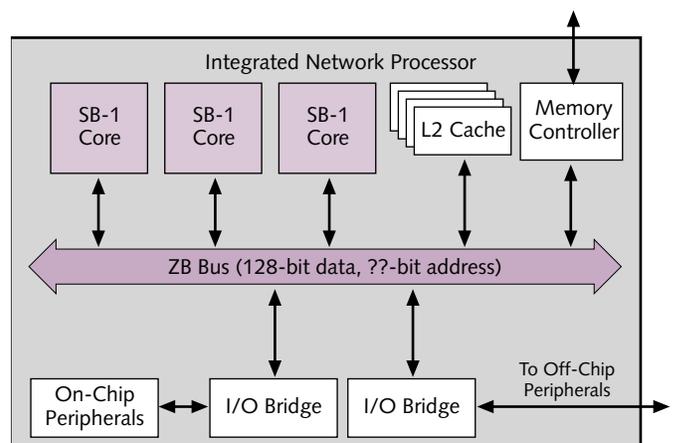


Figure 3. SiByte plans to introduce NPUs that integrate multiple SB-1 cores on a die with memory and peripherals. The practical number of cores that can be integrated is limited mainly by manufacturing considerations.

Price & Availability

SiByte will not license the SB-1. Instead, the core will be the foundation for a family of processors the company will announce in 2H00 and ship in 2001. For more information, go to www.sibyte.com.

the marketplace. However, much depends on how well SiByte's future NPUs provide a total system solution. Intelligent hardware integration, software, and technical support are as important in this market as raw performance is.

Why SiByte Bet On MIPS

As a startup with no historical baggage, SiByte could have based its first design on almost any CPU architecture, or even created its own proprietary alternative. But after examining all the options, SiByte decided the MIPS architecture was the only practical choice.

The company ruled out a proprietary solution, because it's difficult to establish a new architecture in the marketplace. SiByte would have had to convince other companies to port their operating systems, development tools, and application software to the new architecture or write everything from scratch. Besides, a new architecture would probably end up looking very MIPS-like, anyway, as C-Port and Sitera discovered. Working independently, both companies started with the basic MIPS-I architecture, dropped a few instructions, added a few instructions, and arrived at unofficial (unlicensed) MIPS-like architectures for their NPU cores.

Feature	SiByte SB-1	MIPS 20Kc	MIPS 5Kc	Lexra NetVortex	IBM PPC 440
Architecture	MIPS64	MIPS64	MIPS64	MIPS-I*	PowerPC
Architectural Width	64 bits	64 bits	64 bits	32 bits	32 bits
Superscalar Issue	4-way	2-way	No	No	2-way
Issue Order	In order	In order	In order	In order	Out of order
ALU Pipelines	2 x 64 bits	2 x 64 bits	1 x 64 bits	1 x 32 bits	2 x 32 bits
FPU Pipelines	2 x 64 bits	1 x 64 bits	None	None	None
L1 Cache (I/D)	32K/32K	32K/32K	0-64K/0-64K	0-64K/0-64K	0-64K/0-64K
Synthesizable?	No	No	Yes	Yes	No
IC Process	0.15μ	0.15μ	0.15μ	0.15μ	0.18μ
Die Size (w/ Caches)	25mm ²	24mm ²	Varies	Varies	Varies
Core Voltage	1.2V	1.5V	1.5-1.8V [†]	1.5V	1.8V
Core Frequency	1GHz	750MHz	250-275MHz [†]	450MHz	555MHz
Dhrystone 2.1	2,000 mips	1,500 mips	300-350 mips [†]	450 mips	1,000 mips
Power (typical)	2.5W	2W	500mW [†]	n/a	1.4W
Performance Ratio	2 mips/MHz	2 mips/MHz	1.3 mips/MHz [†]	1 mips/MHz	1.8 mips/MHz
Power Ratio	800 mips/W	750 mips/W	650 mips/W [†]	n/a	714 mips/W
Availability	2001	1Q01	1H00	1Q01	Now

Table 1. SiByte's SB-1 has the most powerful microarchitecture of any MIPS-compatible embedded-processor core, but the MIPS64 20Kc might match its performance if the 20Kc could reach an equivalent clock frequency. Also, this table doesn't do justice to Lexra's NetVortex, which has several features (such as rapid context switching) that boost performance beyond that indicated by Dhrystone mips scores. (*Lexra's cores are not fully MIPS compatible. [†]Based on estimates by MIPS Technologies. n/a= not available.)

Many networking customers are partial to MIPS. General-purpose CPUs based on MIPS workstation/server cores are popular in network equipment, and development tools are plentiful. That's why Lexra is pursuing the market with its MIPS-like NetVortex core. But SiByte didn't want to take the chance of being sued by MIPS Technologies, which filed a patent lawsuit against Lexra last year (see [MPR 12/6/99-03](#), "MIPS vs. Lexra: Definitely Not Aligned"). So SiByte coughed up the money—probably a few million dollars—for a genuine MIPS64 license.

ARM has more customers and market share than MIPS has, but ARM's 32-bit uniscalar cores are optimized for low power consumption and aren't nearly powerful enough for the NPUs SiByte envisions. Licensable (and highly configurable) 32-bit cores are also available from ARC Cores and Tensilica, but their architectures are proprietary and not as well supported by tools, operating systems, and software, so SiByte ruled them out. PowerPC has good tool support and high performance, but IBM and Motorola aren't currently licensing PowerPC to other semiconductor companies, and there are no 64-bit embedded PowerPC cores.

That left the MIPS architecture by process of elimination. But none of the off-the-shelf CPU cores on the MIPS roadmap is powerful enough for the NPUs that SiByte had in mind. So instead of licensing the 64-bit 5Kc or 20Kc cores from MIPS, SiByte licensed the MIPS64 ISA and designed its own 64-bit core. It is now the proud owner of the world's most powerful MIPS-compatible embedded-processor core—although the 20Kc is still the most powerful MIPS-compatible core available for general licensing and could match the SB-1's performance if it could reach a comparable clock frequency. (SiByte has no plans to license the SB-1 to anyone else.)

Not every NPU vendor has followed SiByte's line of reasoning. The NPUs from IBM, Intel, and Agere (see [MPR 6/12/00-02](#), "Agere's Pipelined Dream Chip") are based on mostly new and proprietary architectures—but then, IBM and Intel are huge companies that swing much more weight in the marketplace than SiByte, and Agere was recently acquired by Lucent. C-Port and Sitera elected to go with simpler MIPS-like 32-bit cores and make up the difference by integrating several of them on a chip with lots of supporting logic. Lexra, still fighting a rearguard action against MIPS, chose a similar solution, based on its R3000-class core. HyperChip and Chameleon

(see *MPR 6/12/00-01*, “Chameleon Crosses CPU, FPGA”) have announced NPUs based on ARC’s configurable cores.

It may seem risky for SiByte to enter the fray against so much competition. But the SB-1 provides strong evidence that SiByte’s experienced engineering team is nobody’s door-mat. If the SB-1 lives up to expectations, and if the resulting NPUs are as well designed as the core, SiByte has a bright future—and is a prime acquisition target. 

To subscribe to Microprocessor Report, phone 408.328.3900 or visit www.MDRonline.com