BROADCOM SHOWS OFF NEW CPU

BRCM 5000 Provides Dual Issue, Multithreading for Broadband Chips

By Linley Gwennap {11/22/10-01}

....

One should not light a candle and hide it under a bushel, but rather place it where all can see its light. For years, Broadcom has developed its own CPUs and hidden them inside a variety of system-on-a-chip (SoC) products. Recently, the company provided documentation on these CPUs exclusively to *Microprocessor Report,* allowing us to reveal its accomplishments for all to see. These documents provide an in-depth portrayal of an entire family of CPUs, culminating in the most recent design, the BRCM 5000 (code-named Zephyr).

The 5000 is a superscalar MIPS32 CPU that supports two threads for increased performance. The design uses a 12-stage pipeline to achieve a production speed of 1.3GHz. It supports a dual-level cache subsystem and can be combined into a dual-CPU configuration for a total of four threads. This level of performance is well suited to emerging set-top boxes (STBs) and other high-end consumer products. The initial 5000-based products, including the BCM7420 processor, entered production in 65nm LP in 3Q09. The first 40nm BRCM 5000 products entered production in 4Q10.

Broadcom has four major business units: enterprise, broadband, connectivity, and mobile. The latter two rely on ARM CPUs, because of their focus on low power and the large base of ARM software in the mobile market. For these products, the company licenses ARM CPU cores. For enterprise and broadband, however, Broadcom has relied extensively on MIPS CPUs. Although it has licensed some MIPS cores, the company has a long history of developing its own CPUs and today uses these CPUs for all of its broadband products, including DSL, cable-modem, fiber, STB, and digital-TV (DTV) chips. (Some of these products include MIPS-designed CPUs as well.) These broadband products generate more than $2 billion in annual revenue.

Generations of CPU Development

Broadcom's CPU development stretches back more than 10 years, when the company shipped its first DSL chip with an internally designed CPU, now known as the BRCM 3300 core. This CPU, which the company still uses today, is quite small (less than 1mm² in 40nm CMOS, including cache), making it ideal for low-cost products such as ADSL chips. The simple scalar CPU uses a six-stage pipeline to achieve clock speeds of 700MHz in 40nm G, as Figure 1 shows. It generally uses 32KB of instruction cache and 16KB of data cache, although the synthesizable design can be configured in multiple ways. The small data cache is optimized for broadband applications, in which most data is streamed through the processor and doesn't need to be cached.

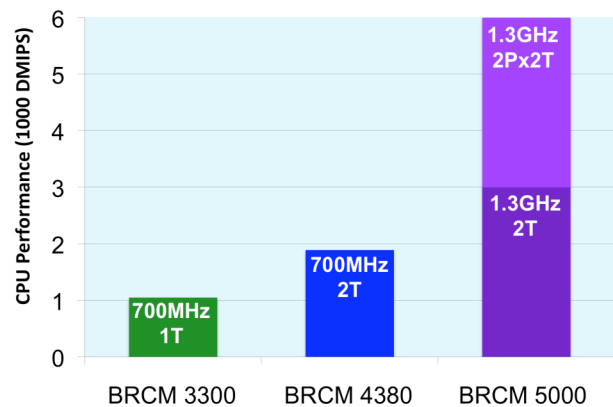Broadcom later developed the BRCM 4355 CPU for higher-performance modems (e.g., VDSL) and set-top

Figure 1. Three generations of Broadcom CPUs. The company has developed several MIPS CPUs, all of which it continues to use in various consumer SoCs. 2T=two threads; 2P=two CPUs. (Source: Broadcom)

boxes. The 4355 is a dual-CPU module in which the two CPUs share a single data cache. This structure was originally designed to simultaneously support two operating systems—for example, an RTOS to process broadband data and Linux to provide a user interface. The shared 32KB data cache simplifies synchronization and avoids wasting cache on streaming broadband data. Each of the two CPUs is similar to the 3300 in speed and issue rate.

A more recent version of this CPU, called the BRCM 4380, expands the shared data cache to 64KB and adds an exclusive 128KB level-two cache. The 4380 also adds DSP extensions and a floating-point unit to accelerate 3D user interfaces in set-top boxes. This CPU, which achieves a speed of 700MHz in 40nm G, appears in Broadcom products such as the BCM7400, BCM7405, and BCM7335 STB processors.

The company started a parallel CPU effort in 2000 when it acquired SiByte, a startup that had developed its own high-end MIPS CPU (see *MPR 6/26/00-04*, "SiByte Reveals 64-Bit Core for NPUs"). By 2005, however, Broadcom decided that integrated SoCs were a better investment than standalone processors, so it merged the SiByte CPU team with its in-house efforts. The new team brought extensive expertise in high-performance CPU design, allowing Broadcom to develop a much more powerful CPU: the BRCM 5000. This CPU implements the MIPS32 v2 instruction set with Broadcom-specific DSP and multithreading extensions. In addition to the BCM742x series of STB processors, the 5000 appears in cable-modem processors and DTV processors.

BRCM 5000's Super Pipeline

To design the BRCM 5000's pipeline, Broadcom initially took the six-stage pipeline from its earlier CPUs and broke each stage into two, a classic example of superpipelining. As Figure 2 shows, the team then adjusted some of the stages to optimize the timing. The pipeline starts with the usual fetch and decode stages. Each cycle, the CPU fetches four instructions from the 32KB instruction cache. To support multithreading, the fetch unit alternates fetching from each of the two threads. If a branch is detected in the decode stage (N3), the CPU accesses the branch predictor and redirects

BRCM 3300 Pipeline

BRCM 5000 Pipeline

Figure 2. Comparison of 6-stage BRCM 3300 pipeline to 12-stage BRCM 5000 pipeline. The latter allows extra stages for cache accesses and adds stages for instruction buffering (Buf) and dual instruction issue (Iss).

the next instruction fetch. A taken branch thus results in a two-cycle pipeline bubble, but because the other thread uses one of those cycles, the affected thread sees only a single-cycle bubble.

The new stage N4 is a buffering stage that covers up these taken-branch bubbles. Because the fetch unit handles twice as many instructions per cycle (IPC) as the execution unit, it will typically get ahead of the instruction execution, causing the buffer to fill. When a taken-branch bubble occurs, the execution unit can continue executing from the instruction buffer without stalling. The instruction buffer has 64 entries: 32 for each thread.

To support the 5000's dual-issue capability, the new stage E0 examines the two instructions at the front of the instruction buffer to make sure that they can be executed together (e.g., they have no dependencies) and issues them to their appropriate execution units. Instructions are always issued and executed in program order. If the execution pipeline stalls at any point, the issue unit will stop issuing instructions, causing the instruction buffer to fill until the stall is cleared.

The 5000 has two integer execution units. Both can handle any simple integer operation. One unit also handles shift, integer multiply and divide, and DSP instructions. The second handles load/store instructions. Despite the superpipelining, the ALU requires only a single cycle, eliminating any latency from back-to-back ALU operations. Accessing the data cache takes an additional two cycles, however, resulting in a three-cycle load latency. The compiler will attempt to insert additional instructions between the load and the use instructions to avoid stalling the pipeline.

Conditional branches are not resolved until stage E3, so the branch-misprediction penalty is 11 to 14 cycles. To avoid this penalty, the 5000 includes sophisticated branch prediction, including a 4K-entry Gshare-based branch history table (BHT), a 64-entry branch target buffer (BTB), and an 8-entry return address stack (RAS). All three of these structures are doubled for the two threads.

The design also includes a 64-bit FPU that executes floating-point add and multiply instructions in six cycles. The unit is fully pipelined, so it can start a new add or multiply each cycle. Although this floating-point performance is modest compared with that of PC processors, it is well suited to the basic 3D interfaces and graphics that are often used in set-top boxes.

Multithreading Adds Little Area

Instead of the 4355's dual-CPU model, the 5000 implements two threads per CPU. Like other multithreaded designs, the 5000 appears to software as two virtual CPUs, each with its own register files and user state. In order to achieve this appearance, the CPU duplicates these
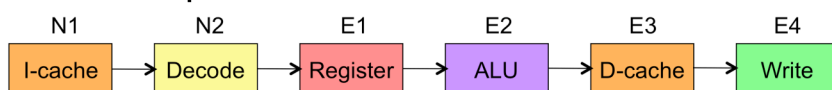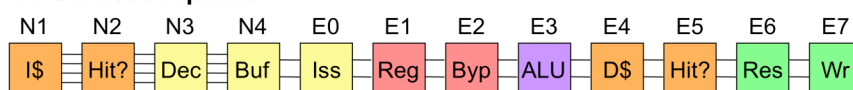
registers in hardware, allowing instructions from both threads to execute simultaneously. In many cases, the two instructions that are issued together will be from different threads, giving the CPU an effective execution rate of one instruction per cycle per thread. But when one thread stalls, perhaps during a memory access, the other thread can use both issue slots, increasing its throughput to two instructions per cycle.

Other than duplicating the register files, multithreading adds little complexity to the overall design. Each register address is extended by 1 bit to contain the thread number. Dependency checking is the same once the thread number is taken into account—instructions from different threads are never dependent. Instructions flow through the pipeline in the same way regardless of which thread they are from. The instruction buffer is divided in two to separate the instruction streams; that way, if one thread is stalled, the CPU can continue to issue instructions from the other thread. Broadcom estimates that adding the second thread increased the die area by only 10%.

Figure 3 shows Broadcom's simulated performance figures for four similar CPU designs, showing the effect of both dual issue and dual threading. In a single-issue design, the second thread increases performance by as much as 60% to 70% for applications with poor cache-hit rates. In these applications, the IPC of a single thread is 0.3 to 0.4, so the second thread has plenty of empty instruction slots to fill. For typical applications with a good cache-hit rate and an IPC of about 0.6, however, the improvement of the second thread drops to about 25%.

The BRCM 5000, however, is a dual-issue design, doubling the number of instruction slots. As a result, the second thread provides a gain of 60% to 70% across a much broader range of cache-hit rates. We expect most applications to fall in this range, although applications with a single-thread IPC above 1.0 will naturally see some degradation in performance for the second thread.

This comparison does not take into account the decrease in cache-hit rate caused by sharing the caches between two threads. When the second thread is enabled, the 32KB instruction and data caches are essentially reduced to 16KB per thread, although sharing data and dynamically allocating the cache between the threads yields some benefits. Accounting for these issues, Broadcom estimates that the net performance gain for the second thread is about 30%. Still, this gain is much greater than the modest increase in die area. In addition, applications with minimal cache misses do better: on CoreMark 1.0, the CPU achieves 1.83 per megahertz using a single thread and 2.99 per megahertz using two threads—a 63% improvement for the second thread.

Caches Minimize Pipeline Disruption

The 5000's instruction and data caches are 32KB each and are four-way set associative. Both caches are parity pro-

tected. The data cache has a 32-byte line size, but the instruction cache has a longer 64-byte line, because instruction accesses are more likely to be sequential. The CPU predecodes instructions as they are loaded into the I-cache, which contains an extra 3 bits per instruction (32 bits) to hold the pre-decode flags. This information speeds the decoding process during the instruction execution.

The primary caches connect to a level-two (L2) cache through a 128-bit interface. The size of the L2 cache can vary but is typically 256KB in Broadcom's current products. The L2 cache is eight-way associative and has a 128-byte line. Unlike the 4380's L2 cache, the 5000's is inclusive, meaning it includes a copy of the data in the primary caches. Although this approach reduces the L2 cache's effective capacity, it simplifies coherence in a multicore design, because the primary caches don't need to be checked. To save power and simplify timing, the L2 cache operates at half the CPU speed. The processor can detect 1-bit and 2-bit errors in the L2 cache data.

Like most MIPS-compatible designs, the 5000 implements a joint TLB that translates both instruction and data addresses. Because this TLB is visible to software, the CPU includes two copies, one for each thread. Each copy contains 64 entries. The design also includes smaller instruction and data TLBs that facilitate single-cycle access to the primary caches. The ITLB and DTLB each have 32 normal (4KB) page entries. The ITLB has one entry for variable-size pages; the DTLB has four such entries. These TLBs are duplicated for each thread.

One Design for Five Fabs

Like most CPUs, the BRCM 5000 uses extensive clock gating to turn off dormant pipeline stages and registers on a cycle-by-cycle basis. Software can also set the CPU's clock divisor to 2, 4, 8, or 16, or it can use a WAIT instruction to turn off the clock to the entire CPU. The design does not, however, implement voltage gating.
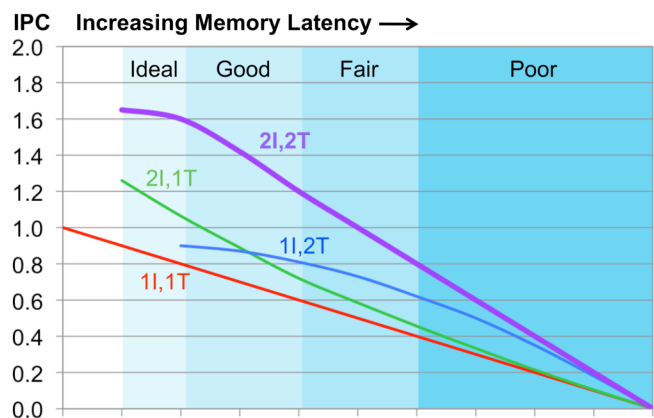

Figure 3. Performance improvement for dual-issue (2I) and dual-threaded (2T) CPU designs. Dual threading provides a sizable performance gain, particularly when coupled with a dual-issue design. (Source: Broadcom)

	Broadcom BRCM 5000	ARM Cortex-A9	ARM Cortex-A9	Intel Atom CPU	MIPS 74K
Design Type	Hard core	Hard (Osprey)	Soft core	Hard core	Soft core
Instruction Set	MIPS32	ARMv7	ARMv7	x86-64	MIPS32
CPU Speed*	1.3GHz	1.3GHz‡	1.0GHz‡	1.6GHz	1.2GHz
Issue Rate	2 per cycle	2 per cycle	2 per cycle	2 per cycle	2 per cycle
Instr Reordering?	—	Yes	Yes	—	Yes
Multithreaded?	2 threads	—	—	2 threads	—
Pipeline Depth	12 stages	9 stages	9 stages	16 stages	15 stages
L1 Cache (I/D)	32KB/32KB	32KB/32KB	32KB/32KB	32KB/24KB	32KB/32KB
Special Instructions	DSP, FPU	DSP, FPU, Neon, Jazelle	DSP, FPU, Neon, Jazelle	SSE3, FPU	DSP, FPU
CoreMark 1.0	3,890 CM	3,745 CM	2,880 CM	5,105 CM	3,050 CM
CM 1.0/MHz	2.99/MHz	2.88/MHz	2.88/MHz	3.19/MHz	2.55/MHz
CPU Die Area†	$1.9mm^2$	$3.4mm^2$	$1.35mm^2$	$6mm^2$‡	$1.5mm^2$‡
CPU Power†	0.37mW/MHz	0.4mW/MHz‡	0.4mW/MHz‡	0.56mW/MHz‡	0.43mW/MHz
IC Process	40nm G	40nm G	40nm G	45nm high-k	40nm G
Chip Production	4Q10	2Q11	2Q10	2Q08	2011 (40nm)

Table 1. Comparison of Broadcom's BRCM 5000 CPU with ARM, MIPS, and Intel CPUs. *Production-silicon clock rate, no overvoltage; †includes L1 caches as shown. (Source: vendors, except ‡The Linley Group estimate)

Unlike earlier Broadcom CPUs, the 5000 uses a combination of custom circuit design and synthesized logic. The CPU team developed custom circuits for critical data-path logic, such as the single-cycle ALU. The company's Central Engineering group provided custom memory blocks such as SRAM, content addressable memory (CAM) for the TLBs, and dual-port register files. The floor plan was hand generated to minimize wire delays between blocks, and the clock tree was manually tuned to control skew.

Broadcom has an unusual manufacturing strategy for its chips. Instead of working closely with a single foundry, the company has created a unified set of design rules that allows it to tape out to five foundries at once: TSMC, UMC, Chartered (now GlobalFoundries), SMIC, and Silterra. This approach prevents Broadcom from taking advantage of a particular foundry's unique features that could boost performance, but it allows the company to more easily manage capacity issues and to aggressively negotiate wafer pricing. In other words, this strategy gives up some clock speed to reduce cost.

To maintain a low manufacturing cost, Broadcom's designs use six metal layers, as opposed to eight in some state-of-the-art processes. (Each additional metal layer can add 10% in cost.) The 40nm BRCM 5000 uses a conservative 10-track library with no LVT (low threshold voltage) transistors. With these parameters, the BRCM 5000 CPU, including L1 but not L2 caches, uses $1.9mm^2$ of die area.

The 65nm LP version is in production at 750MHz, which is achieved in the slow-slow process corner with 10% voltage margin at a maximum junction temperature of 125°C. In 40nm G, the CPU is sampling at 1.3GHz under similar conditions. In a fast-fast process, this design could operate in excess of 2GHz. Broadcom has measured the 40nm CPU (including L1 caches) at 0.37mW per megahertz when running Dhrystone.

The newer version of the 5000 includes a technology that Broadcom calls Adaptive Voltage Scaling (AVS). The chip contains certain test circuits that determine if it is operating near the fast-fast corner or the slow-slow corner. These test circuits contain both analog and digital functions to get a precise reading of the transistor characteristics. For a chip with fast, leaky transistors, the supply voltage is internally lowered, reducing both leakage and transistor speed, but the fast transistors can still achieve the rated clock speed even at the lower voltage. Conversely, the voltage is increased for chips with slow transistors, boosting their performance. Thus, AVS reduces the rated worst-case power, which only occurs in fast-fast chips, while improving speed yield. Because both the test circuits and voltage scaling are built entirely onto the die, AVS requires no binning or external circuitry; everything is adjusted automatically at power-on time.

Meeting the Competition

Although Broadcom does not offer its CPU except in its own SoC products, we have chosen to compare it with other CPUs that are often used in consumer SoCs, including ARM's Cortex-A9 and MIPS Technologies' 74K (see *MPR 5/29/07-01,* "MIPS 74K Goes Superscalar"). We have also included Intel's Atom CPU, which the company will use in future consumer SoC products. Cortex-A9 is available in both soft (synthesizable) and hard (custom) versions; both appear in Table 1. Intel's Atom CPU is currently in production in the second-generation Moorestown processor (see *MPR 5/31/10-01,* "Intel Cuts Atom's Power"). These designs are all instantiated in TSMC's 40nm G process except Atom, which uses Intel's 45nm technology with high-k metal gates.

All four of these CPUs can issue two instructions per cycle, but their designers have taken two different approaches to filling those instruction slots: multithreading and instruction reordering. When an instruction stalls, the BRCM 5000 and Atom can continue executing instructions from the other thread. When an instruction stalls in the 74K

or Cortex-A9, the CPU can execute subsequent instructions from the same thread. Instruction reordering is more complex than multithreading, requiring extensive logic to determine which instructions can be reordered, keep track of them during execution, and put them back in order at the end of the pipeline. This logic increases the die area and power of a reordering CPU. In addition, the performance gains from reordering are typically less than the gain from a second thread. At 2.99 CoreMarks per megahertz, the dual-threaded BRCM 5000 scores better than the single-threaded Cortex-A9 (2.88) and 74K (2.55).

At 1.3GHz, the BRCM 5000 is as fast as any of these CPUs except Atom, which uses a deep pipeline and Intel's advanced process technology. Atom is typically sold at 1.6GHz but is available at faster speeds with overvoltage. Cortex-A9 is typically licensed as a soft core; few products use the Osprey hard core, which is rated at 2.0GHz "typical" with overvoltage but 1.6GHz with production margins. Without overvoltage, Osprey is rated at 1.3GHz. The 74K is rated at 1.2GHz without overvoltage in a 12-track library; in the same 10-track library that the 40nm BRCM 5000 uses, it would probably operate at 1.1GHz. As discussed previously, Broadcom uses conservative design rules and a sizable voltage margin to maximize yield and reduce manufacturing cost. Even so, the 5000 is faster than the licensed soft cores that most of Broadcom's competitors use.

The BRCM 5000's die area is less than that of any of the hard cores in this group and slightly more than the synthesizable cores. Intel's Atom CPU, which uses a 64-bit architecture and includes full PC compatibility, is two to three times larger than its RISC competitors. Not coincidentally, Atom also lags in power dissipation. The BRCM 5000, on the other hand, offers the best power per megahertz, and this rating does not include the effects of AVS, which will reduce power in many real-world situations.

As its second-generation BRCM 5000 enters production, Broadcom is hard at work on its next CPU, codenamed Frenzy. This design will use instruction reordering and issue four instructions per cycle, greatly increasing both single-thread and multithread performance. The first processors based on the Frenzy CPU will sample in 2012.

Frenzy should be comparable to ARM's recently announced Cortex-A15 CPU (see *MPR 11/22/10-02,* "Cortex-A15 Eagle Flies the Coop").

Do-It-Yourself Upgrades

Why is Broadcom investing in CPU development when it could simply license CPUs from MIPS? One explanation is simply to reduce its license fees. The company consumes more than 100 million MIPS-compatible CPUs per year, so the savings per core adds up.

More importantly, the BRCM 5000 solves a problem that the 74K does not. Some of Broadcom's large broadband customers need to run two operating systems on the same processor, a task that would require two separate 74K CPUs. Using multithreading, Broadcom's 4355 CPU supports two operating systems with only a slight increase in die area over a single-thread CPU. By the time MIPS introduced the multithreaded 34K CPU (see *MPR 2/27/06-01,* "MIPS Threads the Needle"), the dual-issue 5000 was available with considerably more performance than the single-issue 34K.

In addition, the 5000's performance using two threads is better than that of Cortex-A9 or the 74K, and it uses less power than those designs. Although the differences are modest, they add up, giving Broadcom an edge in performance per watt and performance per dollar (die area). Compared with Intel's Atom, the 5000 is more power efficient and uses less die area. Looking to the future, Broadcom's Frenzy design looks like it will outperform every licensable CPU that MIPS has yet announced.

Developing its own MIPS-compatible CPUs allows Broadcom to better meet the needs of its customers. With a strong CPU design team built from SiByte and other sources, the company has developed a high-performance CPU, the BRCM 5000, that is entering production in 40nm, and it is developing a new CPU that will further boost performance in future products. These designs show that Broadcom is not just an SoC integrator but has a strong CPU capability as well. This capability should help the company fend off competition in its core broadband and consumer-video markets. ♦