Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

MicroBlaze v7 Gets an MMU

Memory Manager Brings Full-Fledged Linux to Xilinx Processor Core

By Tom R. Halfhill {11/13/07-01}

Xilinx is upgrading its MicroBlaze embedded-processor core again, this time adding an optional memory-management unit (MMU) that allows the 32-bit processor to run sophisticated operating systems supporting virtual memory. Developers can also substitute a simpler memory-protection unit (MPU) or omit supervised memory management altogether.

The first full-fledged operating system announced for the new MicroBlaze v7 is Lynuxworks BlueCat Linux. Until now, MicroBlaze processors were limited to simpler embedded operating systems that don't support virtual memory or memory protection. With its optional MMU or MPU, MicroBlaze v7 is suitable for a wider range of embedded applications requiring greater security and reliability.

MicroBlaze v7 has other improvements as well. New instructions provide faster floating-point performance and better I/O with coprocessors and custom logic. In addition, Xilinx has upgraded the CoreConnect interface to the latest CoreConnect Processor Local Bus (PLB) v4.6 specification, which provides faster links to on-chip peripherals.

All these improvements strengthen MicroBlaze's position against competing processor cores intended for synthesis in FPGAs—primarily Altera's Nios II and ARM's new Cortex-M1. Those rival processors don't have MMUs or MPUs. MicroBlaze v7 is available now as part of a $495 development kit. (An extra $100 buys a development board with a Xilinx Spartan-3E FPGA.) That price includes a MicroBlaze v7 license, and developers owe no royalties when deploying the processor in Xilinx chips.

Robust Memory Management

Xilinx has steadily improved MicroBlaze since introducing the soft processor in 2001. Two years ago, Xilinx began offering an optional FPU. (See *MPR 5/17/05-02*, "MicroBlaze Can Float.") Last year, Xilinx lengthened the instruction pipeline to permit higher clock speeds. (See *MPR 11/13/06-01*, "Xilinx Revs Up MicroBlaze.") And earlier this year, Xilinx released MicroBlaze v6, which added a few other minor enhancements. Now, with MicroBlaze v7, Xilinx is introducing big-league memory management, which significantly expands the range of embedded applications for which MicroBlaze is suited.

With an MMU, MicroBlaze can run full-fledged operating systems based on the Linux 2.6 kernel. Windows CE is another possibility, although Xilinx hasn't announced a deal with Microsoft yet. Processors that support virtual memory can run powerful operating systems, larger application programs, and multiple programs while avoiding memory collisions that would compromise security and reliability. In addition, virtual memory allows a system to work with less physical memory, which can reduce costs and power consumption. As small embedded systems increasingly tackle heavy-duty applications once limited to bigger systems, sophisticated memory management is becoming almost a necessity.

Of course, many embedded systems don't need this level of memory management, so the MicroBlaze MMU is optional. It's just another configuration option available when MicroBlaze developers synthesize the processor using the Xilinx development tools. Another option is to implement an MPU, which provides memory protection without virtual

Memory Manager Type	LUTs (Virtex-5)	LUTs (Spartan-3)
Memory Management Unit (MMU)	910	1,100
Memory Protection Unit (MPU)	560	670
Privlieged Mode Only	34	38

Table 1. Sizes of three MicroBlaze v7 memory-management options. Each option occupies additional lookup tables (LUT) in the FPGA, enlarging the processor core. (A fully configured core requires about 3,000 LUTs.) Note that the sizes for these options are slightly different in Xilinx Virtex-5 and Spartan-3 FPGAs, because the LUTs are different: Virtex-5 LUTs have six inputs, whereas Spartan-3 LUTs have four inputs.

memory and address translation. An MPU is appropriate for embedded systems that must shield a program's memory region against accidental or malicious intrusions by other programs. Yet another option is to implement privileged-mode execution without memory protection or virtual memory. In privileged mode, only the operating system or a privileged application program can execute certain instructions critical to system security.

Table 1 shows how each option (MMU, MPU, or privileged execution) affects the size of the synthesized processor, as measured by the number of additional lookup tables (LUT) required to implement the feature in the FPGA's programmable-logic fabric. Implementing privileged mode alone requires so few LUTs that it's practically a no-brainer. The other options demand more forethought. In particular, the MMU—which requires about 1,000 LUTs—will account for approximately one-third of a fully configured MicroBlaze v7 core. (To put that size in perspective, the Spartan -3E 1600E chip on a MicroBlaze v7 development board has about 33,000 LUTs and costs $10 to $15, depending on volume.)

The full-blown MMU is the largest memory-management option, partly because it needs a translation lookaside buffer (TLB) to cache a subset of the table that translates virtual and physical memory addresses. MicroBlaze v7 has a 64-entry unified TLB. To supplement this software-managed buffer, there are shadow entries for instruction-memory pages and data-memory pages. The number of shadows is user configurable: one, two, four, or eight entries for instructions and the same for data. (The default configuration is two shadows for instructions and four shadows for data.) The processor automatically manages the shadows, which prevent the TLB from thrashing. Memory pages can range in size from 1KB to 16MB, and mixed sizes are allowed. With 32-bit effective addressing, MicroBlaze v7 can address up to 4GB of flat memory.

The MicroBlaze MMU is patterned after the one in an IBM Power 405 processor. That's no coincidence. Some Virtex-family FPGAs integrate a hardened Power 405 core, which is much faster than a MicroBlaze processor synthesized in the fabric. Having a similar MMU brings a few advantages to MicroBlaze v7. First, programmers will have an easier time porting virtual-memory operating systems from the Power Architecture to MicroBlaze. Second, developers should find it easier to create a multicore design that mates one or more MicroBlaze cores with a Power 405 in a shared-memory configuration. And third, the Power-style MMU prepares MicroBlaze v7 for possible future arrangements with IBM to integrate newer Power cores into Xilinx FPGAs.

Faster CoreConnect Bus

CoreConnect is IBM's on-chip bus for SoCs, introduced in 1999. (See *MPR 7/12/99-03*, "PowerPC 405GP Has CoreConnect Bus.") Although IBM created CoreConnect mainly for its own Power Architecture processors, anyone can freely license CoreConnect as synthesizable intellectual property (IP), and it's not specific to a particular CPU architecture. Over the past eight years, soft-IP vendors have made many of their licensable peripheral cores compatible with CoreConnect. The only on-chip bus supported more widely is ARM's AMBA.

For greater efficiency, CoreConnect separates low-speed and high-speed peripherals on separate buses joined by a bridge. Until now, MicroBlaze supported only the slower On-chip Peripheral Bus (OPB), which has a 32-bit-wide datapath. MicroBlaze v7 still supports the OPB for backward compatibility but adds support for the faster Processor Local Bus (PLB). The PLB's datapaths are configurable during synthesis for 32-, 64-, or 128-bit widths. Bus bandwidth depends on the width of these datapaths and the FPGA's clock frequency, which can reach 550MHz in the fastest Virtex-5 devices. At 550MHz, the maximum theoretical bandwidth of a 128-bit PLB would be 8.8GB/s.

The PLB connects directly to the CPU and provides a multidrop bus shared by several on-chip peripherals. It supplements a proprietary Xilinx interface called the Fast Simplex Link (FSL). An FSL is a direct point-to-point interface, not a multidrop bus. FSLs are faster than shared buses but require more logic gates per I/O interface. An SoC design can use one or more FSLs in combination with a CoreConnect PLB for different purposes, giving developers a wealth of options.

Figure 1 shows an example of an SoC implemented in a Xilinx FPGA with a MicroBlaze or Power 405 processor core. In this example—a TCP/IP packet processor—the most critical datapaths link the Ethernet controller to the external-memory controller and the CPU. Those datapaths are FSLs, which can be 32 to 128 bits wide. Less critical components share a CoreConnect PLB. The Xilinx development tools can automatically configure an FSL for a particular purpose or allow developers to configure the interface manually.

With IBM's blessings, Xilinx has slightly modified the standard CoreConnect IP. These modifications were necessary because the programmable-logic gates of FPGAs aren't as efficient as the standard-cell gates of ASICs, especially when routing signals across a large chip. Datapaths and clock trees tend to stretch much longer in FPGAs, making timing closure more difficult. The problem gets worse in complex designs that distribute numerous peripherals throughout the chip on a shared bus. Consequently, Xilinx has modified the PLB to be more synchronous and to eliminate indeterminate

Price & Availability

MicroBlaze v7 is available now as part of Xilinx Embedded Development Kit 9.2, for $495. The kit includes MicroBlaze configuration tools, Eclipse-based software-development tools, and other software, plus documentation. For $595, the kit comes with a development board, a Spartan-3E 1600E FPGA, and a JTAG probe. Xilinx doesn't charge royalties for deploying MicroBlaze designs in Xilinx FPGAs. For more information, see:

- *www.xilinx.com/microblaze*

data bursts. Xilinx says these changes, though relatively minor, allow developers to attach 10 or 20 peripherals to the CoreConnect PLB without timing problems.

In addition, Xilinx has modified the PLB to work more efficiently with hardened transceivers built into some Virtex-5 FPGAs. These transceivers are a PCI Express endpoint and a trimode Ethernet media-access controller (TEMAC), which deliver much better performance than would equivalent soft-IP controllers synthesized in the fabric. The PCI Express endpoint is fully buffered and supports 1, 2, 4, or 8 lanes. The TEMAC transceiver supports 10Mb/s, 100Mb/s, and Gigabit Ethernet rates. In a Xilinx benchmark test, a packet processor based on MicroBlaze v7 and the TEMAC achieved 750Mb/s raw throughput—an impressive 75% utilization of the transceiver's maximum theoretical bandwidth.

Although AMBA enjoys wider support than CoreConnect does, the latter standard is better for MicroBlaze. The Power 405 processor built into some Virtex-family FPGAs uses CoreConnect, so it's easier for developers to create asymmetric multiprocessors around the Power 405 and MicroBlaze cores. Many peripheral-IP cores from third-party vendors work with either AMBA or CoreConnect, usually by adding a simple gasket adapter.

New Instructions Boost Performance

Xilinx has added 11 new instructions to MicroBlaze v7: three for floating-point operations, and eight for use with FSLs. The new floating-point instructions are straightforward. One instruction, FSQRT, calculates the square root of a 32-bit floating-point value in either 27 or 29 clock cycles, depending on whether the MicroBlaze processor is configured with a three- or five-stage pipeline. (The deeper pipeline is faster.) Without using FSQRT, the same operation performed by a function call to a software library would take about 500 cycles.

The other two new floating-point instructions convert integers to floats or vice versa. The FLT instruction converts a 32-bit integer into a 32-bit float in four or six clock cycles, depending on the pipeline depth. Calling the same function in software would take 330 cycles. Conversely, the FINT instruction converts a 32-bit float into a 32-bit integer in five or seven clock cycles, depending on the pipeline depth. A software function call would take 88 cycles.

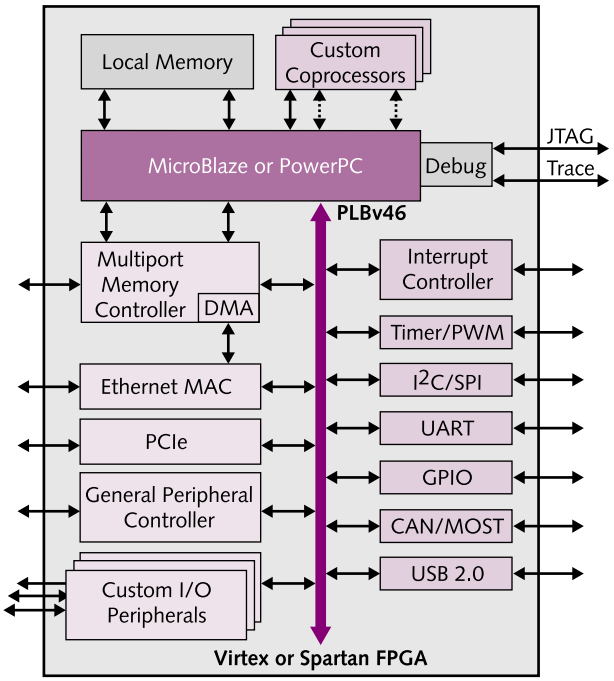

Figure 1. Example SoC block diagram. This packet processor uses Xilinx Fast Simplex Links (FSL) for critical datapaths and a shared IBM CoreConnect Processor Local Bus (PLB) for other on-chip peripherals. MicroBlaze v7 is the first version of the core to support the CoreConnect PLB; previous MicroBlaze cores supported only the slower CoreConnect On-chip Peripheral Bus (OPB). The multiport memory controller is special Xilinx IP with built-in DMA; it's compatible with DDR1 and DDR2 SDRAM.

Eight new instructions improve I/O when a coprocessor is connected to a MicroBlaze core via an FSL. These instructions take the form of PUT and GET operations, and they allow programs to manage the coprocessor I/O as blocking or nonblocking operations. In a blocking relationship, the CPU stops processing other operations until it handles a coprocessor's request for attention. In a nonblocking relationship, the CPU continues processing other operations while the coprocessor's requests enter a FIFO buffer. The CPU isn't blocked unless the buffer fills. Developers can configure the size of the buffer according to the coprocessor's needs.

In addition, MicroBlaze v7 can accommodate twice as many FSL interfaces as before (16 vs. 8), and programs can dynamically assign coprocessors to individual FSL interfaces at run time. Previously, coprocessor assignments to FSLs were hard-coded into the user's software. Any changes required developers to recompile the software. Dynamic assignments allow developers to write software that adapts to changing conditions and workloads. For example, a developer could create precompiled software libraries that select their appropriate coprocessors at run time, according to the custom

Feature	MicroBlaze v6 Embedded Dev Kit 9.1	MicroBlaze v7 Embedded Dev Kit 9.2
TCP/IP Stack	LwIP	Treck
On-Chip Interconnect	CoreConnect On-chip Peripheral Bus (OPB)	CoreConnect Processor Local Bus (PLB v4.6)
External Memory Controller	Xilinx MCH_OPB_DDR	Xilinx Multiport Memory Controller
Ethernet MAC	Fast Ethernet (10–100Mb/s)	Gigabit Ethernet Trimode MAC (10–100–1,000Mb/s)
Xilinx FPGA	Virtex-5 LX	Virtex-5 LXT
Packet Throughput	~70Mb/s	>250Mb/s

Table 2. MicroBlaze v6 versus MicroBlaze v7 performance. This comparison is based on a packet-processor design pitting the new MicroBlaze core against the older version. The upgraded design is more than three times faster. However, Xilinx also changed other variables, including the TCP/IP stack, memory controller, and Ethernet controller, which clouds the comparison somewhat. In particular, the network interface leaps from a soft-IP Fast Ethernet MAC (100Mb/s) to a hard-wired Gigabit Ethernet Trimode MAC (10–100–1,000Mb/s).

hardware in the coprocessors and the tasks to be performed. Multimedia-acceleration libraries could run on coprocessors specializing in fast Fourier transforms (FFT) or finite impulse-response (FIR) filters.

Table 2 shows the results of porting a packet-processor design from MicroBlaze v6 to v7—throughput improved more than 3x, from about 70Mb/s to more than 250Mb/s. However, notice that this comparison (conducted by Xilinx) doesn't isolate the effect of each variable changed in the design. In particular, the Ethernet controller in the upgraded design is much faster. Nevertheless, the comparison demonstrates what is possible. Increasing the maximum theoretical bandwidth of a system doesn't guarantee higher throughput, and one feature of MicroBlaze v7 is better CoreConnect support for the TEMAC hard-wired into Virtex-5 LXT chips.

New Kid on the Block: ARM

MicroBlaze v7 is the second new version of the processor that Xilinx has introduced this year, and the third since 2006. Although these steps have been incremental, they add up to a significantly better processor. The quickening pace of improvement may be due to the arrival of fresh competition: ARM's Cortex-M1.

The Cortex-M1 is the first ARM processor core sanctioned for deployment in FPGAs and optimized for their programmable-logic fabrics. Previously, ARM allowed licensees to test their designs in FPGAs but not to deliver finished designs in the chips. ARM's course change was prompted partly by the rising costs of designing and manufacturing ASICs, and partly by the popularity of the Xilinx MicroBlaze and Altera Nios II cores. (Xilinx and Altera have sold tens of thousands of licenses for their processors.) The Cortex-M1 is a major new development that alters the competitive landscape. (See *MPR 3/19/07-01*, "ARM Blesses FPGAs.")

The first FPGA vendor to announce support for the Cortex-M1 was Actel, a much smaller company than Altera or Xilinx. Actel has a special arrangement with ARM to sell Cortex-M1 FPGAs without requiring customers to purchase an ARM license or pay chip royalties to ARM. This deal dramatically lowers the cost of deploying an ARM-based design. Xilinx hasn't announced a similar arrangement yet, but *MPR* suspects it's a possibility in the future. Although the Cortex-M1 and MicroBlaze processors would seem to make rivals of ARM and Xilinx, their relationship remains more cooperative than competitive. ARM understands that MicroBlaze is primarily a loss-leader product that Xilinx created to sell more FPGAs. A MicroBlaze v7 license costs only $495, so the chips—not the licenses—are the real moneymakers. Xilinx is happy to see customers buying its FPGAs to use with the Cortex-M1, too.

Even so, while ARM and Xilinx cordially shake hands, MicroBlaze v7 is slapping the Cortex-M1 upside the head. The brand-new ARM processor suffers in comparison with the Xilinx loss leader. Although MicroBlaze v7 is priced at a pittance, it's embarrassingly rich in features missing from the Cortex-M1, such as an optional FPU, MMU/MPU, 32-bit divider, and instruction/data caches. On top of that, MicroBlaze can reach higher clock frequencies than the Cortex-M1 can. ARM's biggest selling point is that the Cortex-M1 is from ARM. The ARM architecture is almost an industry standard, and it's supported by tons of peripheral IP, development tools, and software.

As Table 3 shows, Altera's Nios II is a closer match for MicroBlaze, even though it hasn't been significantly upgraded since 2004. (See *MPR 6/28/04-02*, "Altera's New CPU for FPGAs.") The addition of an optional MMU/MPU gives MicroBlaze v7 its first big advantage over Nios II. However, Altera retains one advantage: a user-configurable instruction-set architecture. Nios II developers can create custom instructions to accelerate specific applications, which can dramatically improve performance. To achieve similar results, MicroBlaze developers can implement coprocessors in the programmable-logic fabric. (Coinciding with the online publication of this article on November 13, Altera is announcing an arrangement with Synopsys that will make it easier for developers to move Nios II-based designs from FPGAs to standard-cell ASICs. *MPR* plans to cover this development in the future.)

Note that the price gaps among these processors are shrinking. Before the Cortex-M1, the difference between licensing a processor core from an FPGA vendor or licensing one from ARM was four orders of magnitude: about $500 for a MicroBlaze or Nios II versus millions of dollars for an ARM. With the Cortex-M1, ARM has departed from its long-standing practice of not publicly disclosing licensing fees. Although the exact price for a Cortex-M1 license is still under wraps, ARM says it will cut deals for less than $100,000—a huge price break. And, as described above, Actel sells preconfigured Cortex-M1 FPGAs without requiring an

Feature	Xilinx MicroBlaze v7.0	Xilinx MicroBlaze v6.0	Altera Nios II/f	Altera Nios II/s	Altera Nios II/e	ARM Cortex-M1
Architecture	MicroBlaze	MicroBlaze	Nios II	Nios II	Nios II	ARMv6-M
Primary FPGA Targets	Virtex-5 Spartan 3	Virtex-5 Spartan 3	Stratix, Cyclone, HardCopy	Stratix, Cyclone, HardCopy	Stratix, Cyclone, HardCopy	Fusion, ProASIC-3, Stratix, Virtex-4/5, Cyclone, Spartan
Config. ISA	—	—	Yes	Yes	Yes	—
Pipeline Depth	3 or 5 stages	3 or 5 stages	6 stages	5 stages	1 stage*	3 stages
I-Cache	0–64K	0–64K	0–64K	0–64K	—	—
D-Cache	0–64K	0–64K	0–64K	0–64K	—	—
Local Memory	0 or 2 256K each	0 or 2 256K each	0–8 Configurable	0–4 Configurable	—	0 or 2 1K–1,024K each
32-Bit Multiplier	Optional	Optional	Optional	Optional	—	Yes, 2 options
32-Bit Divider	Optional	Optional	Optional	Optional	—	—
Barrel Shifter	Optional	Optional	Optional	Optional	—	Yes
FPU	Optional 32 bits (New instructions)	Optional 32 bits	Optional 32 bits	Optional 32 bits	Optional 32 bits	—
Branch Predict	—	—	Dynamic	Static	—	—
Priv. Levels	1 or 2	1	2	2	2	1
Coprocessor Interface	FSL (New instructions)	FSL	Avalon-MM	Avalon-MM	Avalon-MM	AMBA-Lite
On-Chip Interconnect	CoreConnect PLB v4.6	CoreConnect OPB	Avalon-MM	Avalon-MM	Avalon-MM	AMBA-Lite
Memory Management	Optional MPU or MMU	—	—	—	—	—
Translation Lookaside Buffer (TLB)	Optional 8-entry I + D 64-entry unified	—	—	—	—	—
Core Freq (Max)	220MHz †	220MHz †	205MHz	165MHz	200MHz	Up to 72MHz ‡ >170MHz**
Int. Perf (Max)	240 Dmips	240 Dmips	225 Dmips	127 Dmips	31 Dmips	0.8 Dmips / MHz
FP Perf (Max)	50 MFLOPS	50 MFLOPS	n/a	n/a	n/a	n/a
FPGA Logic Cells	980–3,000	960–1,700	1,800	1,150	600	4,300+ LUT3 tiles (~1,900 LUT4 cells)
Introduction	Nov 2007	Mar 2007	2004	2004	2004	4Q07
Price	$495	$495	$495	$495	$495	<$100,000 (ARM) Free (Actel)

Table 3. Feature comparison of the Xilinx MicroBlaze v7, MicroBlaze v6, Altera Nios II, and ARM Cortex-M1 processor cores. All are 32-bit embedded processors designed and optimized for synthesis in FPGAs. Altera's Nios II processor is available in three basic configurations that developers can customize. Key differences between the MicroBlaze v7 and v6 are highlighted in purple text. With its new optional MMU/MPU, MicroBlaze v7 gains an important advantage over Nios II; the Cortex-M1 fares poorly in this comparison but is redeemed by a more familiar CPU architecture. (FSL is the Xilinx Fast Simplex Link. Avalon-MM is Altera's memory-mapped interface.) *Nios II/e has a six-stage pipeline, but it works like a one-stage pipe. †Maximum clock speed assumes synthesis in the fastest Xilinx Virtex-5 FPGAs. ‡Estimate for synthesis in an Actel ProASIC3 or Fusion FPGA. **Estimate for synthesis in a Xilinx Virtex-5 FPGA. (n/a: not applicable)

ARM license at all. With Altera and Xilinx practically giving away their FPGA-ready processors, ARM had to modify its own licensing model to make the Cortex-M1 competitive.

The Future of CPUs for FPGAs

As FPGA prices fall, and the costs of ASICs rise, deploying an SoC in programmable logic becomes increasingly attractive. As MPR has noted in the past, the production volume at which an FPGA implementation becomes more economical than an ASIC implementation keeps tilting in favor of FPGAs, and we perceive nothing on the horizon that will alter that trend. For that reason, the future of MicroBlaze (and Nios II) looks bright.

However, one thing that could change is the CPU architecture that developers choose to implement in an FPGA. For now, MicroBlaze and Nios II are by far the most popular choices, because they are promoted by their respective FPGA vendors and are practically free. ARM's Cortex-M1 changes the equation by introducing the option of using the industry's most popular 32-bit embedded-processor architecture. In time, other embedded-processor architectures may join the fray. At present, ARC, MIPS Technologies, and Tensilica don't necessarily forbid licensees from deploying designs in FPGAs, but they don't encourage it, either—and they have not optimized their processors for programmable logic.

If other CPU architectures do become available for FPGAs, at affordable prices, the time may come when the FPGA vendors' own CPU architectures look less attractive. Although Altera and Xilinx have sold many more CPU

licenses than even ARM ever will, a great number of those licenses are purchased by students and tinkerers who never intend to mass-produce their designs. Companies seriously intending to produce FPGA-based SoCs in volume may prefer to use a more widely supported CPU architecture. MicroBlaze and Nios II could become interesting footnotes in the history of microprocessors.

Even if that happens, the Altera and Xilinx processors will have served their purposes. They are selling more FPGAs, they are helping to seed the market for FPGA-based SoCs, and they are defining the features and optimizations that FPGA-specific processors should have. Whether or not MicroBlaze and Nios II live long and prosper, they are wise investments for their vendors.