

ARM's DIGITAL SIGNAL CONTROLLER

New Cortex-M4 Brings DSP Extensions to Cortex-M Family

By Tom R. Halfhill {4/12/10-01}

ARM is pitching its new Cortex-M4 processor as a digital signal controller (DSC)—the first time ARM has so described one of its processor cores. Essentially, a DSC crosses a digital signal processor (DSP) with a microcontroller (MCU) for double duty in controller

applications that need a little signal processing. But, in fact, the Cortex-M4 is not a departure for ARM. It's more like a bridge between the ARM9, ARM11, and Cortex-M families.

The Cortex-M4 has no new integer DSP instructions that veteran ARM programmers haven't seen before. Instead, it adopts the same DSP and SIMD extensions introduced with the ARM9E processor core in 1999, later inherited by the ARM11 family in 2002. (See [MPR 6/21/99-03](#), "Arm Refocuses DSP Effort.") Other ARM processors sharing the same DSP instructions are the Cortex-R4, Cortex-A5, Cortex-A8, and Cortex-A9.

In essence, the Cortex-M4 provides a Cortex upgrade path for existing ARM9 and ARM11 designs. It can also upgrade designs based on fellow members of the Cortex-M family—the Cortex-M0 and Cortex-M3. For ARM, the Cortex-M4 adds a modicum of signal processing to the Cortex-M product line while nudging customers toward a newer core that reaps higher licensing fees.

It's not entirely a marketing ploy, though. The Cortex-M4 improves performance by guaranteeing single-cycle throughput for DSP operations that formerly required multiple clock cycles. It also has a new, higher-performance, optional FPU, should an application require single-precision floating-point arithmetic. (Some control applications do.)

Lastly, the Cortex-M4 offers programmers better compatibility with software-development tools across the whole Cortex family, unlike the older ARM9 and ARM11 tool

chains. ARM says five companies have already licensed the new processor.

Minimal ISA Saves Memory

Code-named Merlin, the Cortex-M4 builds on the foundation of the Cortex-M3, which made its debut with the ARMv7 architecture in 2004. (See [MPR 11/29/04-01](#), "ARM Debuts Logical V7.") The Cortex-M3 remains in the product line and is positioned for MCUs that don't need the Cortex-M4's DSP extensions. Two years ago, ARM enhanced the Cortex-M3, adding new power-saving features and compatibility with third-party fault-tolerant technology. (See [MPR 5/12/08-01](#), "Fault Tolerance for Cortex-M3.")

Both processors are simple 32-bit cores with three-stage pipelines, like ARM's best-selling (but aging) ARM7TDMI. Both processors omit caches, whose nondeterministic behavior is a liability in hard real-time systems. Both processors have a modified Harvard bus architecture that provides separate AHB-Lite 32-bit I/O interfaces for instructions and data. As Figure 1 shows, the Cortex-M4 also has a separate AHB-Lite interface to the rest of the system.

All processors in ARM's Cortex-M series adhere to the ARMv7-M instruction-set architecture (ISA), a stripped-down subset of the full ARMv7 ISA. Notably, these processors cannot execute the vast majority of traditional 32-bit ARM instructions. The Cortex-M0 and Cortex-M1 are the smallest members of the ARMv7-M family, supporting only a few 32-bit instructions and several 16-bit Thumb and




Figure 1. Cortex-M4 block diagram. The new processor has much in common with the six-year-old Cortex-M3. Key differences are the DSP extensions inherited from the older ARM9E and ARM11 families and an optional single-precision FPU. ARM refers to the floating-point version of the processor as the Cortex-M4F. Optional blocks are light purple in this diagram.

Thumb-2 instructions. (See [MPR 3/2/09-01](#), “ARM’s Smallest Thumb,” and [MPR 3/19/07-01](#), “ARM Blesses FPGAs.”)

In addition to supporting Thumb and Thumb-2, the Cortex-M3 has nearly a hundred 32-bit instructions. The Cortex-M4 goes further, supporting all the same 16- and 32-bit instructions as the Cortex-M0, Cortex-M1, and Cortex-M3, plus eighty 32-bit DSP/SIMD instructions. (With the optional FPU, 29 more 32-bit instructions join the party.) ARM refers to the Cortex-M4’s extended ARMv7-M ISA as ARMv7-ME. These overlapping ISAs are confusing to describe, but Figure 2 clarifies their relationships.

ARM says the Cortex-M4 is designed for software development in C, not assembly language, although programmers often end up writing critical routines in assembly code. To ease their labor, ARM will soon introduce a new library of DSP routines for the Cortex-M4. These routines will resemble the libraries offered by Analog Devices, ARC International (now part of Virage Logic), Ceva, Tensilica, Texas Instruments, and other CPU/DSP vendors.

Single-Cycle DSP Performance

ARM says its C compiler will automatically use some DSP instructions without programmer intervention. Most instructions will require programmers to invoke intrinsic functions. Although it’s still “C programming,” it’s only a step above assembly language. High-level C compilers for DSPs have come a long way in recent years, but they’re still

not as good as a skilled assembly-language programmer, and ARM has less experience in this regard than other companies focusing more exclusively on signal processing.

As mentioned above, the Cortex-M4 improves on the ARM9E and ARM11 families by converting their multi-cycle DSP operations into single-cycle operations. The instruction mnemonics and operations remain the same—they’re just faster. The additional logic required for these faster implementations becomes insignificant when the processor is fabricated in modern process technology.

Table 1 shows the 16- and 32-bit multiply-accumulate (MAC) operations that now execute at a throughput rate of one instruction per cycle. These include 16-bit dual MACs, which the Cortex-M4 can sustain at a single-cycle rate.

Some limitations apply. The Cortex-M4 lacks caches and the single-cycle X/Y data memories commonly found in full-fledged DSPs. It also lacks a special interface for closely coupled memory, though designers can attach zero-wait-state memory to the AHB interface. The Cortex-M4’s simple three-stage pipeline cannot execute load/store instructions in parallel with math operations, but the processor can pack two 16-bit values in one 32-bit load or store.

Other limitations: the Cortex-M4 doesn’t support circular addressing, bit-reverse addressing, or zero-overhead looping—distinguishing features of real DSPs. (As partial compensation, it can minimize the overhead of loops by prefetching branch addresses.) There are no wide (e.g., 40-bit) accumulators or other provisions for overflows.

Without dedicated accumulators, the Cortex-M4 must store the results of 32-bit operations in general-purpose 32-bit registers. Some operations generate 64-bit results, which are stored in pairs of 32-bit registers. Results of dual 16-bit operations are packed into a single 32-bit register. All DSP instructions operate on 8-, 16-, or 32-bit operands. To crunch 24-bit data, programmers must use 32-bit operations and discard the low eight bits.

Without dedicated accumulators, the Cortex-M4 must store the results of 32-bit operations in general-purpose 32-bit registers. Some operations generate 64-bit results, which are stored in pairs of 32-bit registers. Results of dual 16-bit operations are packed into a single 32-bit register. All DSP instructions operate on 8-, 16-, or 32-bit operands. To crunch 24-bit data, programmers must use 32-bit operations and discard the low eight bits.

It’s plain from these limitations that the Cortex-M4 isn’t a true DSP. ARM doesn’t claim it is. Even for a DSC, it’s minimally equipped. However, the Cortex-M4 will be suitable for control applications that need a little signal processing. ARM pitches the Cortex-M4 for motor controllers, power management, industrial control, automobiles, and audio systems.

When audio codecs written for the Cortex-M3 are optimized for the Cortex-M4, ARM says the new processor can

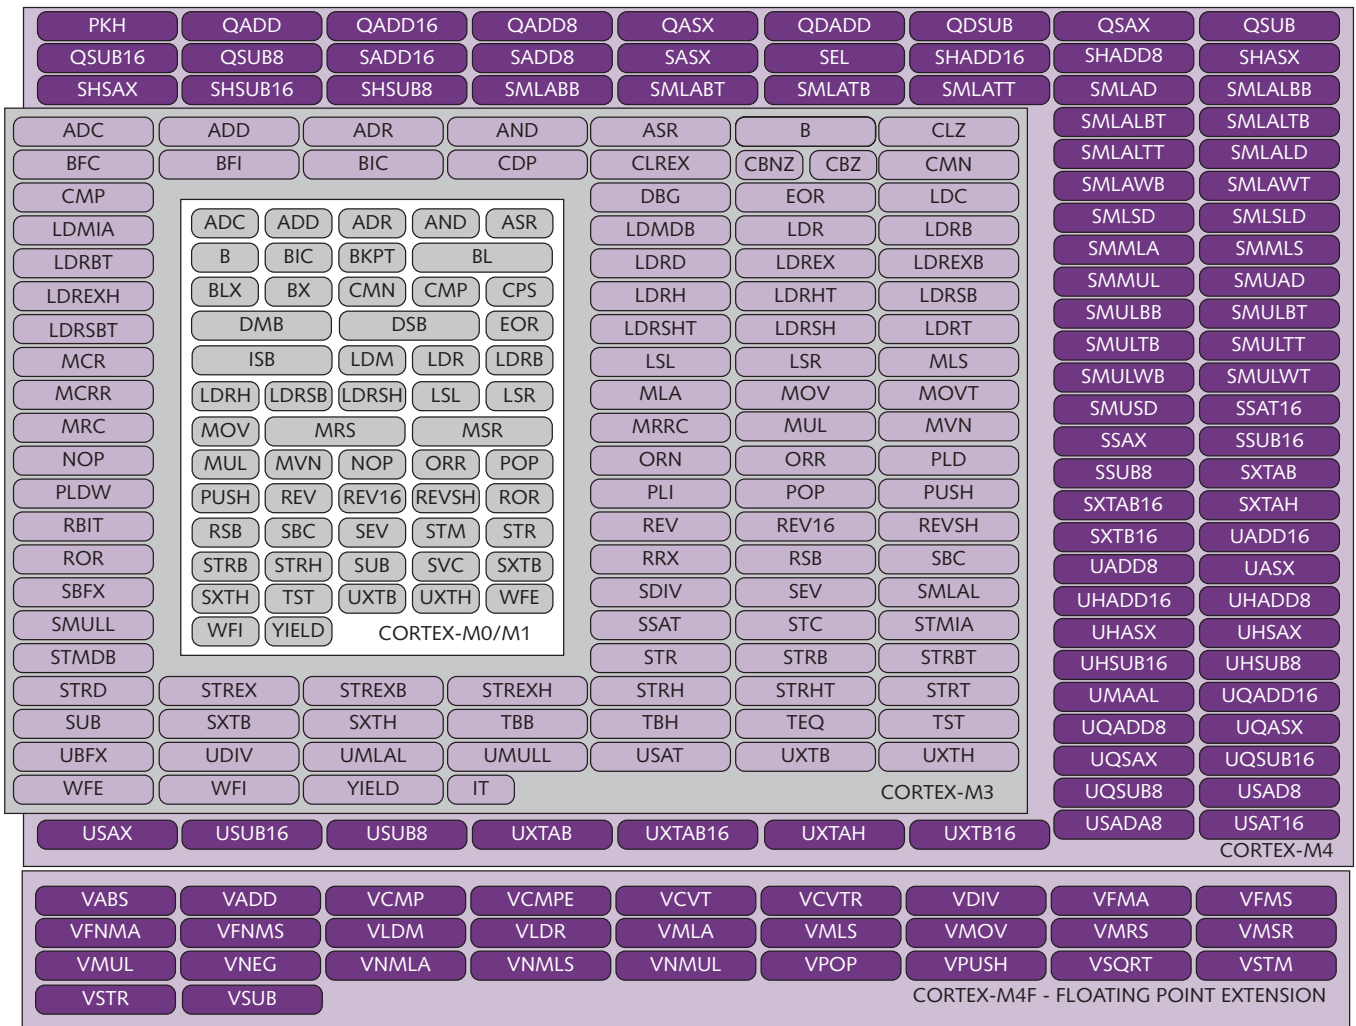


Figure 2. The ARMv7-M ISA as implemented in the Cortex-M family. None of these processors supports the older 32-bit ARM instruction set. Instead, they support subsets of 32-bit instructions and 16-bit Thumb and Thumb-2 instructions. Their compact ISAs improve code density in memory-constrained embedded systems. The Cortex-M4 has the largest number of 32-bit instructions in the Cortex-M family, thanks to the addition of DSP/SIMD instructions and optional floating-point instructions.

decode high-quality 320Kb/s MP3 files at about 9.5MHz while consuming only 0.5mW. That’s much better than an MCU without DSP extensions. Even so, it falls a little short of 32-bit CPU/DSP cores from Tensilica and Virage Logic that are optimized for audio.

Optional FPU Adds Registers

With its optional single-precision FPU, the Cortex-M4 becomes the Cortex-M4F. The FPU is an entirely new design for the Cortex-M4 and adds 25,000 gates to the 65,000-gate integer core. As Figure 2 showed, the FPU comes with 29 new 32-bit instructions. Otherwise, the two cores are identical.

Some control systems require floating-point math, so FPUs are no longer strangers to MCUs. The FPU in the Cortex-M4F is respectably fast, performing most operations in one to three clock cycles. However, most instructions take an additional cycle if the following instruction needs the

result. Table 2 lists the floating-point instructions and their execution latencies.

The FPU has its own three-stage instruction pipeline that’s decoupled from the three-stage integer pipeline. In some cases, integer instructions can complete out-of-order while long-latency floating-point instructions are executing. The longest-latency floating-point instructions are 32-bit divisions and square roots, which have a nominal latency of 14 clock cycles. These operations can execute in only one cycle when not followed by other floating-point instructions.

Unlike most other 32-bit processors, ARM processors have only 16 general-purpose registers. This limitation would cause register congestion if the FPU had to share them. Instead, the FPU is generously provisioned with its own register file—32 registers, 32 bits wide. Pairs of these 32-bit registers can store 64-bit double-word values.

Overall, the FPU is adequate for the Cortex-M4’s intended market. Most other 32-bit processor cores designed for

MAC Operation	Cortex-M4 Instructions
16-Bit Operations	
$16 \times 16 = 32$	SMULBB, SMULBT, SMULTB, SMULTT
$16 \times 16 + 32 = 32$	SMLABB, SMLABT, SMLATB, SMLATT
$16 \times 16 + 64 = 64$	SMLALBB, SMLALBT, SMLALTB, SMLALTT
$16 \times 32 = 32$	SMULWB, SMULWT
$(16 \times 32) + 32 = 32$	SMLAWB, SMLAWT
$(16 \times 16) \pm (16 \times 16) = 32$	SMUAD, SMUADX, SMUSD, SMUSDX
$(16 \times 16) \pm (16 \times 16) + 32 = 32$	SMLAD, SMLADX, SMLSD, SMLSDX
$(16 \times 16) \pm (16 \times 16) + 64 = 64$	SMLALD, SMLALDX, SMLSLD, SMLSLDX
32-Bit Operations	
$32 \times 32 = 32$	MUL
$32 \pm (32 \times 32) = 32$	MLA, MLS
$32 \times 32 = 64$	SMULL, UMULL
$(32 \times 32) + 64 = 64$	SMLAL, UMLAL
$(32 \times 32) + 32 + 32 = 64$	UMAAL
$32 \pm (32 \times 32) = 32$ (upper 32 bits)	SMMLA, SMMLAR, SMMLS, SMMLSR
$(32 \times 32) = 32$ (upper 32 bits)	SMMUL, SMMULR

Table 1. Single-cycle MAC instructions for the ARM Cortex-M4. Most of these instructions require multiple cycles to execute on ARM9E and ARM11 processors. The instruction mnemonics and operations haven't changed, but the Cortex-M4 uses additional logic to reduce execution latencies.

MCUs have optional FPUs too, including some that support double-precision operations. In most cases, double precision is unnecessary for controllers, so ARM's single-precision FPU isn't a serious handicap.

Not a Speed Demon

With its stubby three-stage pipeline, the Cortex-M4 isn't designed for supersonic clock speeds. That doesn't matter because 32-bit MCUs with flash memory generally don't run faster than 100–150MHz, anyway. Otherwise, they would outrun their flash. (Remember, the Cortex-M4 has no caches.)

ARM says the Cortex-M0, Cortex-M3, and Cortex-M4 would all cruise at the same clock frequency if optimized and synthesized for the same fabrication process—more than 300MHz in 65nm or even 90nm technology. In the latest 40nm process, ARM estimates, the Cortex-M4 could exceed 500MHz. In Global Foundries' new low-power 65nm-LPe process—more realistic for an MCU—ARM specifies a maximum worst-case clock frequency of 300MHz. Power consumption is about 0.06mW per megahertz in that process.

For power misers, ARM offers the Cortex-M Low-Power Optimization Package. This package has ARM's own physical intellectual property (IP) for the TSMC CE018FG (0.18-micron) CMOS fabrication process. The IP includes ARM's Artisan Metro 180ULL (ultralow-leakage) cell library and ARM's Power Management Kit. (For more information, see [MPR 3/2/09-01](#), "ARM's Smallest Thumb.")

In any fabrication process, the Cortex-M4 is pretty small. ARM has estimated the size of two different core configurations, both fabricated in Global Foundries' 65nm-LPe process. Both configurations target a clock frequency of 150MHz, which is suitable for an MCU with flash memory. The default configuration has no memory-protection unit (MPU), no FPU, full debug/trace features, 16 interrupts, and no wake-up interrupt controller (WIC). Core size is 0.15mm². Add 0.06mm² for the FPU.

The minimum configuration of the Cortex-M4 has no MPU, no FPU, minimum debug/trace features, only one interrupt, and no WIC. Core size is 0.1mm². That's 33% smaller than the default configuration, but you'd need a microscope to tell the difference.

Standard Parts vs. Licensable Cores

In the DSC market, the Cortex-M4 faces two classes of competitors: standard parts and licensable processor cores. Standard parts include other DSCs, MCUs, and discrete DSPs. In raw performance, virtually all discrete DSPs would stomp the Cortex-M4. But they are a shrinking portion of the signal-processing market, accounting for only about 10% of shipments (by revenue). They are being replaced by DSCs, MCUs, and SoCs that integrate the functions of a CPU and DSP.

Among the standard-part competitors are Analog Devices' Blackfin, Atmel's AVR32, Freescale's 56000 series, Microchip's dsPIC, and TI's C2000 series. Interestingly, all these companies, except Microchip, happen to be ARM licensees. So, to some extent, ARM's modest entry into the DSC market competes with some of its largest customers. It's possible that some of those customers will license the Cortex-M4 for future designs, if the ARM core doesn't overlap their existing products.

In any event, standard-part DSCs are generally better signal processors than the Cortex-M4. They have more features common to full-fledged DSPs, such as zero-overhead looping, flexible addressing modes, wide accumulators, and tightly coupled data memories. Their software-development tools are more geared toward signal processing and they already have libraries of common math routines. The closest standard-part competitor for the Cortex-M4 is probably Microchip's dsPIC, a 16-bit controller with similar DSP capabilities.

Competing CPU/DSP Cores

Of course, the advantage of standard parts is that they're available off the shelf, whereas licensing the Cortex-M4 is only the first step in a development project that will probably take 12–18 months before a finished chip arrives from the foundry.

Most customers in the market for a DSC probably won't regard the Cortex-M4 as direct competition for a standard part. Instead, the most likely Cortex-M4 customers are MCU vendors that will license the core, build it into a product line of standard-part DSCs, and then sell those chips downstream to *their* customers—as they are doing now with the Cortex-M0 and Cortex-M3.

ARM says it has licensed the Cortex-M4 to five companies. Three are known: NXP Semiconductors, STMicroelectronics, and TI. Some MCUs built on the Cortex-M4 are scheduled to hit the market late this year, probably in Q4. More will appear next year.

Other companies may license the Cortex-M4 to design custom SoCs, not MCUs. In those cases, the most direct competitors for the new ARM core are other licensable processor cores with signal-processing capabilities. Almost all 32-bit embedded-processor cores have some features for signal processing, and some have powerful extensions that make them truly formidable CPU/DSPs. We'll narrow our comparison to a few small cores that are suitable for controllers and other deeply embedded applications.

TeakLite-II: One Billion Sold

Ceva is the leading vendor of licensable DSP cores. In terms of core size and power consumption, Ceva's TeakLite-II is the closest match for the Cortex-M4. Of the processors covered in this article, the TeakLite-II is by far the most popular. Ceva has sold more than 50 licenses to companies like Broadcom, Infineon, NXP, Samsung, Sony, and Zoran. It's found in cellphones, home audio systems, storage controllers, medical equipment, and other systems. Ceva says licensees have shipped more than one billion units of the TeakLite-II.

Overall, the TeakLite-II is more DSP than MCU. Unlike the 32-bit Cortex-M4, it's based on a 16-bit fixed-point DSP architecture, though it can squeeze 24- and 32-bit signal-processing operations through its 16- and 36-bit datapaths.

The TeakLite-II has traditional DSP features lacking in the Cortex-M4, such as X/Y data memories (configurable

Single-Precision Floating-Point Operation	Cortex-M4 Instruction	Execution Latency (Clock Cycles)
Absolute Value	VABS.F32	1
Addition	VADD.F32	1
Compare	VCMP.F32, VCMPE.F32	1
Convert	VCVT.F32	1
Divide	VDIV.F32	14
Load (multiple doubles)	VLDM.64	1 + 2 x number of doubles
Load (multiple floats)	VLDM.32	1 + number of floats
Load (single double)	VLDR.64	3
Load (single float)	VLDR.32	2
Move (5 variations)	VMOV, VMRS, VMSR	1 or 2
Mul Float	VMUL.F32	1
MAC Float	VMLA.F32	3
Mul, Sub Float	VMLS.F32	3
MAC, Negate Float	VNMLA.F32	3
Mul, Sub, Negate Float	VNMLS.F32	3
Fused MAC Float	VFMA.F32	3
Fused Mul-Sub Float	VFMS.F32	3
Fused MAC, Negate Float	VFNMA.F32	3
Fused Mul, Sub, Negate Float	VFNMS.F32	3
Negate Float	VNEG.F32	1
Negate, Mul Float	VNMUL.F32	1
Pop Double Regs From Stack	VPOP.64	1 + 2 x number of double regs
Pop Float Regs From Stack	VPOP.32	1 + number of regs
Push Double Regs to Stack	VPUSH.64	1 + 2 x number of double regs
Push Float Regs to Stack	VPUSH.32	1 + number of regs
Square Root of Float	VSQRT.F32	14
Store (multiple doubles)	VSTM.64	1 + 2 x number of doubles
Store (multiple floats)	VSTM.32	1 + number of floats
Store (single double)	VSTR.64	3
Store (single float)	VSTR.32	2
Subtract Float	VSUB.F32	1

Table 2. ARM Cortex-M4F floating-point instruction set. ARM counts 25 instructions, but there are actually a few more, if variations are included. These instructions comprise the FPv3-D16 ISA, part of the ARMv7-M ISA. FPv3-D16 complies (mostly) with the IEEE 754-2008 standard; the Cortex-M4 supports IEEE denormals and rounding modes in hardware. The floating-point MACs and fused MACs offer higher precision than the 32-bit integer MACs in the base instruction set.

up to 2MB) and separate program memory (up to 2MB). It can transfer two 16-bit operands in parallel from its data memories in a single clock cycle, which keeps the single-cycle 16-bit MAC well fed. The processor stores results in four 36-bit accumulators, with automatic saturation on overflow. It also has the flexible memory addressing, normalization, and bit-field operations expected of a DSP.

On the other hand, the Cortex-M4 looks more suitable in the controller role. It has an optional MPU for memory management, an optional FPU, and supports more external interrupts. If the target application needs 24- or 32-bit signal processing, the Cortex-M4 has a wealth of single-cycle 32-bit instructions, including 32-bit MACs. Control code will consist almost entirely of 16-bit Thumb and Thumb-2 instructions, so the Cortex-M4's overall code density will be close to the 16-bit TeakLite-II's.

Ceva offers extensions for the TeakLite-II that upgrade the processor into the Xpert-TeakLite-II. These extensions add caches, a DMA controller, an interrupt controller, timers, and better power management. The extended core is designed for heavier-duty signal-processing chores in audio subsystems and mobile embedded applications. Nevertheless, the Cortex-M4 still looks like the better controller core, especially for designs already based on ARM that use a popular real-time operating system (RTOS).

For even more-demanding signal processing, Ceva's third-generation TeakLite-III goes beyond both the TeakLite-II and Cortex-M4. However, the TeakLite-III is intended primarily for cellphones and high-end audio applications, not for deeply embedded controllers.

Multiple Choices From Tensilica

Lately, Tensilica has been repositioning itself as a vendor of "dataplane processor units"—licensable CPU cores with powerful DSP extensions. Actually, these processors are based on Tensilica's 32-bit general-purpose RISC architecture. The extensions make the difference. Configurable processors have always been Tensilica's real specialty, and the company offers numerous products built on two recently enhanced cores: Xtensa LX3 and Xtensa 8. (See [MPR 11/30/09-01](#), "Tensilica Tweaks Xtensa.")

Tensilica's most powerful CPU/DSP core is the ConnX Baseband Engine (BBE), but it's overkill when pitted against the Cortex-M4 in deeply embedded DSC applications. ConnX BBE is intended primarily for cellphones, where Tensilica is struggling for a foothold alongside ARM-based applications processors. NTT DoCoMo recently confirmed that it's using Tensilica's earlier Xtensa LX core in a new SoC designed for Long-Term Evolution (LTE) cellular telephony. (See [MPR 8/10/09-01](#), "Tensilica Plays Baseband.")

Against the Cortex-M4, Tensilica has three smaller processor cores that are likely competitors: the ConnX D2 DSP, the Diamond Standard 106Micro, and the Diamond Standard 212GP.

The ConnX D2 is almost the same size as the Cortex-M4 (70,000 gates vs. 65,000 gates). Built on Tensilica's Xtensa LX3 32-bit RISC architecture, it adds dual 16-bit MACs and 40-bit accumulators. Two-way SIMD instructions improve performance with vector operations. (Tensilica offers a vectorizing C compiler that's compatible with TI C6x DSP and ITU-T intrinsic functions.) ARM's Cortex-M4 is the better choice if the target application requires 32-bit MACs.

Not Quite a Perfect Match

Tensilica's Diamond Standard 106Micro and Diamond Standard 212GP cores are smaller options than the ConnX D2. (Tensilica's Diamond Standard processors are preconfigured cores that retain some configurability; see [MPR 3/29/06-01](#), "Tensilica's Preconfigured Cores.") However, neither processor is a perfect match for the Cortex-M4.

Although the 106Micro is strong competition for ARM's processor as a small controller core, it has no DSP extensions. It's a better competitor for ARM's Cortex-M0 and Cortex-M3.

Tensilica's Diamond Standard 212GP has more in common with the Cortex-M4—it's a small controller core with moderate signal-processing capabilities. However, the 212GP, like the ConnX D2, lacks the 32-bit DSP instructions of ARM's processor. Although the 212GP has 32-bit multipliers and dividers for integer arithmetic, DSP instructions are limited to 16-bit fixed-point operations.

Otherwise, the 212GP looks pretty good against the Cortex-M4. The single-cycle 16-bit MAC unit has four dedicated 32-bit registers and a 40-bit accumulator. It supports zero-overhead looping, normalization, sign extension, and minimum/maximum-value instructions—features missing from the Cortex-M4. Unlike ARM's processor, the 212GP has instruction and data caches, three integrated timers, and 32-bit general-purpose I/O (GPIO) ports that bypass the main system bus. Its deeper pipeline (five stages vs. three) achieves higher clock frequencies, up to 416MHz with a speed-optimized design flow in 65nm-LP. When throttled back to match the Cortex-M4's frequency in a similar process, power consumption will be approximately the same.

For audio applications, Tensilica's HiFi-2 processor—another preconfigured core—will outperform the Cortex-M4. Tensilica designed special audio extensions for the HiFi-2, including MAC units that can perform 24- × 24-bit or 32- × 16-bit dual MAC operations. A related core, the HiFi-EP, has a single 32- × 24-bit MAC unit but achieves higher performance with some algorithms, such as the Digital Theater Systems (DTS) codec.

Virage Logic's ARC 610D

For more than ten years, Tensilica's nemesis was ARC International, another vendor of configurable embedded-processor cores. ARC was acquired last year by Virage Logic, which continues to offer ARC's processors, along with a great deal of related IP for SoC design. (See [MPR 9/14/09-01](#), "Summer Shopping Spree.")

In 2005 and 2006, ARC and Tensilica introduced preconfigured versions of their processors for customers who didn't want to configure the cores themselves. One of those preconfigured cores, the ARC 610D, is direct competition for the ARM Cortex-M4.

The ARC 610D is the smallest preconfigured ARC core with DSP extensions. Like the Cortex-M4, it's designed for controllers and deeply embedded applications. Although

the ARC 610D is preconfigured, customers retain the option of configuring it further, including the addition of more-powerful DSP extensions. (See [MPR 3/14/05-02](#), “ARC’s Pre-configured Cores.”)

Table 3 summarizes the features of the ARM Cortex-M4, Ceva TeakLite-II, Tensilica Diamond Standard 212GP, and Virage Logic ARC 610D. All are licensable embedded-processor cores capable of serving as controllers and signal processors. (Note: We excluded the MIPS Technologies MIPS32 24KE processor from this comparison, despite its DSP extensions. The MIPS 24KE is a larger, faster, more power-hungry core that won’t compete directly with these smaller cores for deeply embedded controllers.)

Off the shelf, the ARC 610D has 16/32-bit MUL and MAC instructions with saturation. Instead of sharing general-purpose registers, DSP instructions have their own registers, including 40-bit accumulators. Separate pipelines allow

DSP instructions to execute in parallel with each other and with ALU instructions—a standout feature for such a small processor. The standard DSP extensions also support zero-overhead looping. Overall, the preconfigured ARC 610D exceeds the signal-processing capabilities of the Cortex-M4.

ARC XY Advanced DSP Extensions

More capabilities are available by adding the ARC XY Advanced DSP Extensions to the ARC 610D. Most notably, these optional extensions add separate X/Y memories for data operands, a dedicated DMA engine for those memories, and special memory-addressing modes. Also included is a library of optimized DSP routines, including functions for fast Fourier transforms (FFT), Viterbi error correction, cyclic redundancy checks (CRC), and 24- × 24-bit MACs.

The X/Y memories are closely coupled to the ARC 610D core and are independent of the closely coupled memories

Feature	ARM Cortex-M4	Ceva TeakLite-II	Tensilica ConnX D2	Tensilica Diamond 212GP	Virage Logic ARC 610D
CPU Architecture	ARMv7-ME 32 bits	Ceva TeakLite-II 16 bits	Xtensa LX3 + D2 32 bits	Xtensa LX3 32 bits	ARCompact 32 bits
Pipeline Depth	3 stages	4 stages	5 stages	5 stages	5 stages
L1 Cache (Instr + Data)	—	—	8KB + 8KB	8KB + 8KB	—
Local RAM (Instr + Data)	—	0–2MB instr. 0–2MB X/Y data	0–128KB instr. 0–128KB data	0–128KB instr. 0–128KB data	1KB–512KB instr. 2KB–16KB data
DSP Extensions	ARM9E / ARM11 DSP + SIMD	Ceva 16-bit fixed-point	Tensilica 16-bit DSP	Tensilica 16-bit DSP	ARC DSP or XY Advanced DSP
MAC Instructions	16 × 16 + 32 → 32 32 × 32 + 64 → 64	16 × 16 + 36 → 36	16 × 16 + 32 → 40	16 × 16 + 32 → 40	16 × 16 + 32 → 32 / 40 24 × 24 + 32 → 32 / 40 32 × 32 + 32 → 40 / 64
Peak MACs	1 per cycle (32 bits)	1 per cycle (16 bits)	2 per cycle (16 bits)	1 per cycle (16 bits)	1 per cycle (32 bits)
16-Bit Instructions	Thumb, Thumb-2	16-bit only	16-bit Xtensa	16-bit Xtensa	16-bit ARCompact
FPU	Optional SP	—	Optional SP	—	Optional SP + DP
Memory Mgmt	Optional MPU	—	—	Optional MPU	Optional MPU
Dhrystone 2.1	1.25Dmips / MHz	n/a	1.4Dmips / MHz	1.38Dmips / MHz	1.3Dmips / MHz
Core Freq (Max)	300MHz (GF 65nm-LPe)	320MHz (65nm-LP)	600MHz (TSMC 65nm-G)	416MHz* (TSMC 65nm-LP)	500MHz (TSMC 90nm-GT)
Core Area (Without FPU)	0.15mm ² (GF 65nm-LPe)	0.1mm ² (65nm-G)	0.19mm ² (TSMC 65nm-G)	0.316mm ² (TSMC 65nm-LP)	0.16mm ² (TSMC 90nm-LP)
Power	<0.04mW / MHz (GF 65nm-LPe)	0.05mW / MHz (65nm-LP)	0.054mW / MHz (TSMC 65nm-G)	0.103mW / MHz (TSMC 65nm-LP)	0.03mW / MHz (TSMC 90nm-LP)
Introduction	Feb 2010	2005	2009	2006	2005

Table 3. Feature summary of the ARM Cortex-M4, Ceva TeakLite-II, Tensilica ConnX D2, Tensilica Diamond Standard 212GP, and Virage Logic ARC 610D. All but the TeakLite-II are general-purpose 32-bit RISC architectures with DSP extensions. TeakLite-II is a 16-bit architecture that leans more heavily toward signal processing than control processing, though it can perform both duties in some applications. All these processors achieve high code density with 16-bit instructions (TeakLite-II exclusively) or a mix of 16-, 24-, and 32-bit instructions. Observe the usual caution when comparing the vendor-supplied estimates for core area and power consumption. ARM’s estimate assumes the default configuration of the Cortex-M4: no MPU, no FPU, full debug/trace, 16 interrupts, and no wake-up interrupt controller (WIC), fabricated in Global Foundries’ 65nm-LPe process at 150MHz. Add about 0.06mm² for the FPU. For a minimal configuration of the Cortex-M4 (no MPU, no FPU, minimum debug/trace, one interrupt, no WIC), ARM estimates 0.1mm². *Tensilica’s specifications for the Diamond 212GP are based on Revision B of the processor in a speed-optimized design flow. (n/a: data not available.)

Price & Availability

ARM's Cortex-M4 digital signal controller is a synthesizable processor core available for licensing now. Options include a single-precision FPU and ARM's Cortex-M Low-Power Optimization Package for the TSMC CE018FG (0.18-micron) CMOS fabrication process. This package includes ARM's Artisan Metro 180ULL (ultralow-leakage) cell library and Power Management Kit. ARM doesn't publicly disclose licensing terms or fees. For more information, visit:

www.arm.com/products/processors/cortex-m/cortex-m4-processor.php

for the ALU. Each bank of X/Y memory can range in size from 1KB to 32KB, and the banks can be single- or dual-ported. With dual-ported X/Y memories, the ARC 610D can fetch two DSP operands in a single clock cycle while storing results from a previous DSP operation. The dedicated DMA engine keeps these loads and stores from interfering with the ALU pipeline.

Special X/Y addressing modes allow programmers to use addresses with variable offsets, modulo addressing, or bit-reverse addressing. Some instructions automatically increment memory addresses after execution. Dedicated address generators offload all address calculations from the ALU. These DSP features go well beyond the ARM Cortex-M4 and the 16-bit DSP extensions in Tensilica's Diamond Standard 212GP. Of course, the price of the ARC XY Advanced DSP Extensions is a larger core and higher power consumption.

As a general-purpose controller, the ARC 610D is equally tough competition for the Cortex-M4. Its five-stage pipeline should propel it to higher clock speeds in any given fabrication process than the Cortex-M4's three-stage pipeline. The ARC 610D has closely coupled memories for ALU operations, and they are configurable: 1KB to 512KB for instructions, 2KB to 16KB for data. The 16/32-bit ARCompact instruction set provides good code density, although the Cortex-M4's preponderance of 16-bit Thumb and Thumb-2 instructions may give ARM a slight edge in this regard. Vendor-supplied power-consumption estimates are notoriously difficult to compare, but the ARC 610D and Cortex-M4 are in the same realm.

All told, the ARC 610D is a particularly strong alternative to the Cortex-M4—both as a controller and as a signal processor. Unfortunately for Virage Logic, ARM has advantages that go beyond the processor cores. For one thing, ARM has a larger ecosystem of third-party development tools and soft IP. Also, nearly all the leading 32-bit MCU vendors are committed to ARM, to the exclusion of almost every other licensable CPU architecture. (Exception: Microchip is using the MIPS32 M4K, M14K, and M14Kc cores; see [MPR 11/16/09-01](#), "MicroMIPS Crams Code.")

A Small Step Forward

When judged purely as a signal processor, the Cortex-M4 isn't a big step beyond the 11-year-old ARM9E, which was really ARM's first DSC. The Cortex-M4 accelerates the signal-processing instructions but adds no new integer DSP instructions. Mainly, the Cortex-M4 provides an upgrade path from existing ARM9/ARM11 designs to the newer Cortex family and ARMv7 architecture.

When judged as a controller, the Cortex-M4 is a bigger improvement over the older cores. It has modern debug/trace features, new sleep modes, an all-new optional FPU, a nested vector interrupt controller (NVIC), and an optional wake-up interrupt controller. These attributes will help make the Cortex-M4 a successful DSC core and a better embedded controller than the ARM9E.

Old-timers may recall that even the ARM9E was not ARM's first attempt at something like a DSC. In truth, the ARM9E was ARM's first *successful* DSC. An earlier attempt was ARM's Piccolo, introduced in 1996. (Not to be confused with TI's present-day Piccolo MCUs, members of the C2000 family.) Piccolo failed, mainly because it was a separate core with its own instruction set, not an integrated CPU/DSP. In 1999, ARM tried again with the ARM9E, introducing the signal-processing extensions that live on today in the Cortex-M4. (See [MPR 6/21/99-03](#), "Arm Refocuses DSP Effort.")

With this history in mind, the Cortex-M4 looks like a worthwhile upgrade from the ARM9 and ARM11, and it opens a migration path toward other, more powerful, members of the Cortex family. (Software written for Cortex-M processors is upward compatible with Cortex-R and Cortex-A processors, but not the reverse, unless programmers limit themselves to Thumb-2 instructions.)

Nevertheless, the Cortex-M4 leaves plenty of room for ARM to make a better DSC core in the future—if ARM is willing to compete with its own licensees that already offer more-capable DSPs. ♦

To subscribe to Microprocessor Report, phone 480.483.4441 or visit www.MPRonline.com