An Independent Analysis of the

CEVA-TeakLite-III Digital Signal Processing Core

By the staff of
Berkeley Design Technology, Inc.

OVERVIEW

The CEVA-TeakLite-III is a 32-bit licensable DSP processor core from CEVA that builds on the legacy of the CEVA-TeakLite, CEVA-TeakLite-II, and CEVA-Teak cores. CEVA-TeakLite-III targets a range of portable and high-definition audio applications, along with VoIP and cellular baseband. BDTI recently completed an independent analysis of the CEVA-TeakLite-III core. In this paper, BDTI presents benchmark results for the CEVA-TeakLite-III that quantify its speed, power efficiency, and area efficiency relative to those of several competitors, and analyzes its strengths and weaknesses.

Introduction

The CEVA-TeakLite-III is a licensable 32-bit DSP processor core from CEVA. It is primarily intended for use in audio applications, though it also targets VoIP and cellular baseband. CEVA-TeakLite-III is the third generation of CEVA's TeakLite architecture, succeeding CEVA-TeakLite and CEVA-TeakLite-II. CEVA-TeakLite-III is the first 32-bit core in the TeakLite family. It is fully synthesizable and, according to CEVA, it can operate at up to 550 MHz in a 65 nm process. CEVA-TeakLite-III family members come with varying combinations of L1 and L2 caches and system interfaces. CEVA has also announced the CEVA-HD-Audio platform, which includes a TeakLite-III core along with a configurable audio subsystem.

CEVA's cores span a range of performance points, and include the high performance, quad-MAC CEVA-X1641, the dual-MAC CEVA-X1622 and CEVA-X1620, the three TeakLite variants mentioned earlier, and an older CEVA-Teak core. According to CEVA, TeakLite family cores have been licensed by more than 50 licensees and shipped in more than a billion devices, mainly for mobile baseband and audio applications. (For simplicity, we will leave off the "CEVA-" prefix in the core names for the remainder of this paper.)

TeakLite-III core characteristics are shown in Table 1.

Process	Clock speed	Core Size (no memory)
65 nm	550 MHz	0.47 mm^2
130 nm	335 MHz	1.14 mm^2

TABLE 1. TeakLite-III core characteristics. The 130 nm data has been certified by BDTI.

TeakLite-III competes with a range of general-purpose DSP and CPU cores from vendors such as VeriSilicon, ARM, and MIPS, and also with application-specific audio solutions, such as Tensilica's 330HiFi audio core and ARC's Sound Subsystems cores. It also competes, indirectly, with other architectures (such as the Texas Instruments TMS320C55x) that are used in chip-level products targeting similar applications.

BDTI recently evaluated the DSP performance of the TeakLite-III core using the BDTI DSP Kernel Benchmarks suite. In this paper, BDTI presents the core's benchmark results and provides an analysis of the TeakLite-III's strengths and weaknesses relative to selected competitors.

Architecture and Instruction Set

TeakLite-III is a 32-bit fixed-point DSP processor core that supports single-cycle 32-bit arithmetic and multiplication operations. The 32-bit fixed-point format is common in general-purpose CPUs but somewhat unusual in DSP processors, which tend to use either 16-bit fixed-point or 32-bit floating-point data. (TeakLite-III's predecessors, TeakLite and TeakLite-II, are both 16-bit fixed-point processors.)

TeakLite-III's support for 32-bit data makes sense in the context of its focus on audio applications, since it enables the processor to efficiently provide the higher numeric fidelity required by many audio algorithms. A few other audio processors also use relatively wide data widths—for example, Tensilica's 330HiFi core uses 24-bit data, while ARC's Sound Subsystems cores use data sizes of up to 32 bits.

The TeakLite-III core uses a Harvard memory architecture with separate buses for instructions and data. The core can load up to 64 bits of data per cycle, and (somewhat unusually) TeakLite-III supports loading of unaligned 32-bit data; this capability is useful in unpacking bit streams, and can also help reduce code size.

As shown in Figure 1, the core includes a "DAAU," or Data Address and Arithmetic Unit, which contains two address generators and a "scalar" unit used for control-oriented operations and complex address computations (such as an "extract" operation that's useful for Huffman decoding).

The core itself does not include data or instruction memory; instead, it includes a data memory controller and a program memory controller that are designed to be attached to on-chip memories of up to 4 GB each.

The data path is called the "CBU," or "Computation and Bit Unit," and includes a register file that includes two 32-bit multiply input registers (each of which can be treated as two 16-bit registers) and two 36 bit product registers; a MAC unit with four 36-bit accumulators; an ALU; a BMU (bit manipulation unit); and hardware accelerators for FFT and Viterbi computations.

Figure 1. CEVA-TeakLite-III Block Diagram. (Figure courtesy of CEVA.)

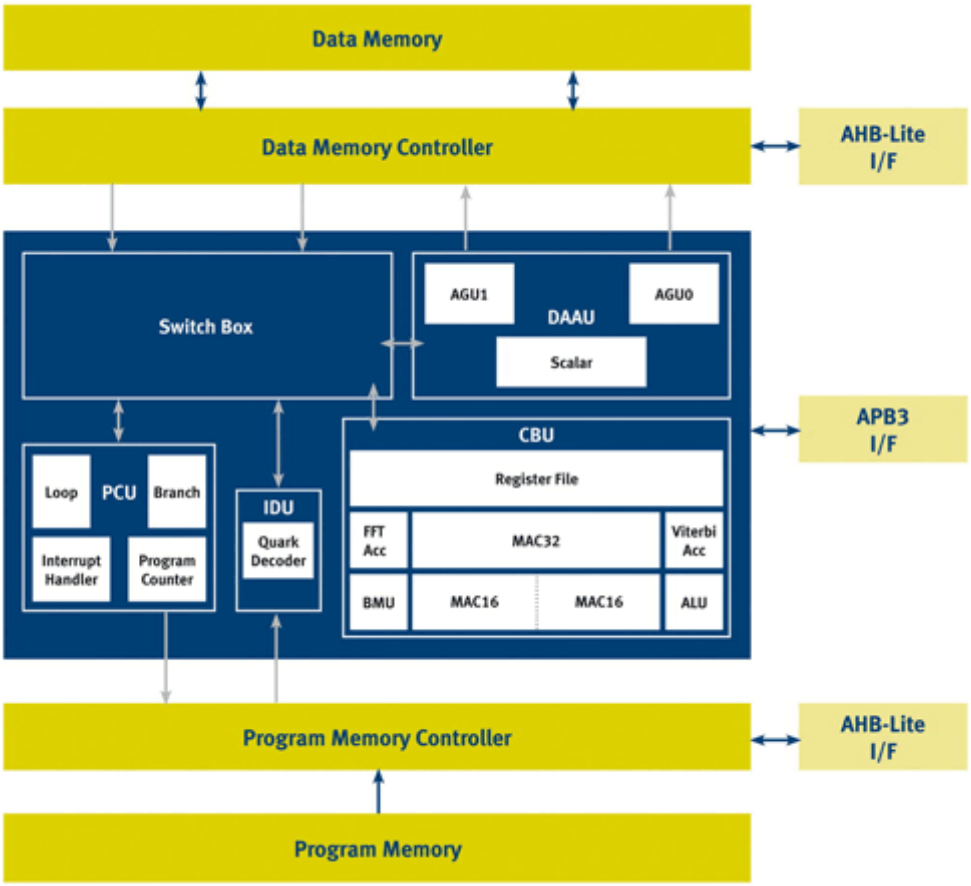

The data addressing unit includes two data address generators and eight data address and arithmetic registers, all of which can be used by the programmer. The data path is able to operate directly on data in memory, rather than requiring data to be loaded to registers first, which can save cycles in some cases. In fact, a few operations can be performed more efficiently if the operands reside in memory rather than in registers or accumulators. For example, TeakLite-III can load a memory operand into an accumulator and swap the lower and higher 16-bit halves in a single-cycle instruction. But swapping the two halves of an operand that already resides in the accumulator requires three instructions and an extra accumulator.

The TeakLite-III MAC unit can execute one 32-bit multiply-accumulate per cycle with 72-bit accumulation (useful for HD audio), or dual 16-bit MACs via SIMD (single instruction, multiple data) instructions. These SIMD MAC capabilities double the per-cycle MAC throughput relative to the TeakLite and TeakLite-II cores. Unlike with some processors, however, on the TeakLite-III there is no way to accumulate two independent MAC results; when two 16-bit MACs are executed in a SIMD fashion, their results are accumulated into a single register.

In addition to its MAC capabilities, TeakLite-III includes a number of specialized features that help accelerate common DSP algorithms. For example, it supports DSP-oriented addressing modes such as modulo and bit-reversed addressing, and supports zero-overhead single- and multi-instruction loops. It also supports fractional multiply modes, along with saturation and rounding. These features help reduce the amount of time (and power) the processor requires for executing signal processing tasks, and they are often missing from general-purpose CPU cores, such as many of the ARM and MIPS cores.

The TeakLite-III instruction set includes specialized instructions for Viterbi decoding and FFTs; these instructions use the accelerators in the CBU. The Viterbi accelerator enables the core to run an ACS (add-compare-select operation) in two cycles, and the FFT accelerator supports two-cycle 16-bit butterfly calculations. A double-precision (32-bit) FFT requires four cycles per butterfly.

TeakLite-III supports a mixed-width instruction set. It has 32-bit instructions, which support multiple parallel operations, and 16-bit compressed "CEVA-Quark" instructions, which allow smaller code size for program segments that don't require top performance. Some instructions (long multiplies, for example) require 48 bits. The mixed-width approach is a common one; for example, ARM has taken a similar approach with its "Thumb" instruction sets. TeakLite-III CEVA-Quark and "regular" instructions can be freely interleaved without requiring a mode change.

Unlike its predecessors, the TeakLite-III core can execute two or three parallel instructions per cycle as part of an "instruction packet." However, the processor is not a VLIW or superscalar architecture; the instruction combinations are fixed and there are a fairly limited number of combinations supported. The processor can, for example, execute a MAC, an add, and a store as a three-way parallel instruction, or a multiply and store as a two-way parallel instruction. Some of the instructions that comprise a parallel instruction can also be executed as standalone instructions, while others are only available within a parallel instruction.

The core includes a program control unit that can fetch up to 64 bits of code at a time in two cycles.

TeakLite-III is backwards compatible with TeakLite, TeakLite-II, and Teak at the assembly source code level, meaning that the newer core can execute assembly code written for the older cores without modification, though using the SIMD capabilities and new instructions will require modification of the existing assembly code.

The TeakLite-III core has a variable ten-stage pipeline. Most instructions use eight stages, while multiplications and multiply-accumulates use ten. This pipeline is much deeper than the four stages used in the TeakLite and TeakLite-II, and enables the core to attain higher clock speeds than its predecessors. To help reduce the number of stall cycles associated with change-of-flow instructions, TeakLite-III supports instruction pre-fetching.

The BDTI DSP Kernel Benchmarks™

The BDTI DSP Kernel Benchmarks are a set of twelve digital signal processing functions that BDTI has independently designed to provide an objective basis for comparing processor performance characteristics—such as speed and memory use—for signal processing applications.

Implementations of the BDTI DSP Kernel Benchmarks functions are carefully optimized to allow a realistic assessment of signal processing performance. BDTI carefully reviews each benchmark implementation in detail to ensure that it complies fully with BDTI's rigorous specifications, and to ensure that the reported performance results are accurate. This verification process is essential to enable fair, apples-to-apples comparisons among processing engines.

BDTI DSP Kernel Benchmarks scores are available for a wide range of licensable DSP cores, packaged DSP

processors, and general-purpose processors. Table 2 lists the twelve BDTI DSP Kernel Benchmarks functions.

Real Block FIR	Two-Biquad IIR	Viterbi Decoder
Single-Sample FIR	Vector Dot Product	Control
Complex Block FIR	Vector Add	256-Point FFT
LMS Adaptive FIR	Vector Maximum	Bit Unpack

TABLE 2. The BDTI DSP Kernel Benchmarks

Benchmark Results

In this section we compare benchmark results for the TeakLite-III core to results for several other licensable cores: CEVA's own CEVA-X1620, ARM's ARM1136, and VeriSilicon's ZSP500. We have also included benchmark results for the Texas Instruments TMS320C55x chip family. (We do not yet have BDTI DSP Kernel Benchmarks for the Tensilica 330HiFi audio core nor the ARC Audio Subsystems cores, and thus cannot make performance comparisons with these cores.)

Though TeakLite-III supports single-cycle 32-bit data computations, the BDTI DSP Kernel Benchmarks are implemented using 16-bit data since this is the data width for which the core attains its highest DSP performance (due to its support for 16-bit SIMD operations). It's important to keep in mind, then, that the performance shown here is not indicative of the core's performance using 32-bit data. All of the other processors shown here also use 16-bit data for their BDTI DSP Kernel Benchmarks implementations.

Speed: BDTImark2000™ and BDTIsimMark2000™

The BDTImark2000 and BDTIsimMark2000 are composite performance metrics that are based on a processor's speed on the full set of BDTI DSP Kernel Benchmarks. BDTImark2000 scores are provided only when a processor's performance has been verified on hardware, whereas BDTIsimMark2000 scores are provided for processors when only simulated results are available. For further information on the BDTI Benchmark suites, the BDTImark2000, and the BDTIsimMark2000, see www.BDTI.com.

The BDTImark2000 and BDTIsimMark2000 are designed to provide a convenient shorthand for processors' signal processing speeds, and are far more accurate than simplified metrics such as MIPS or MFLOPS for this purpose. BDTImark2000 and BDTIsimMark2000 scores for the processors considered in this report are shown in Figure 2.

Comparing performance and efficiency of licensable cores, as opposed to packaged chips, poses some unique challenges. For chips, vendors guarantee that the processor will achieve a certain clock speed. For cores, the clock speed, power consumption, and area depend on the fabrication process, synthesis targets, and other factors. Hence, the performance data (clock speed, power, area) of a core may vary dramatically from one use of that core to the next.

For this reason, BDTI has developed a set of uniform conditions that are used when calculating benchmark results for licensable cores.These conditions ensure that comparisons of benchmark scores at a particular process node use the same fabrication process, and that the reported speed, area, and power metrics for all cores are obtained in the same way. BDTI certifies core performance data that meets these uniform conditions.

For consistency, BDTI calculates scores for licensable cores using projected worst-case data in a uniform process. (Current results use a 0.13 um process; future results will use 90 nm and 65 nm processes.) In this context, "worst-case clock speed" means the clock speed projected for a core assuming worst-case process, voltage, and temperature variations. For packaged processors, scores are computed using the fastest available family member.

During the benchmarking process, the core vendor provides BDTI with clock speed, area, and power results

Figure 2. Overall Speed: BDTIsimMark2000™ and BDTImark2000™ (Higher is Faster)

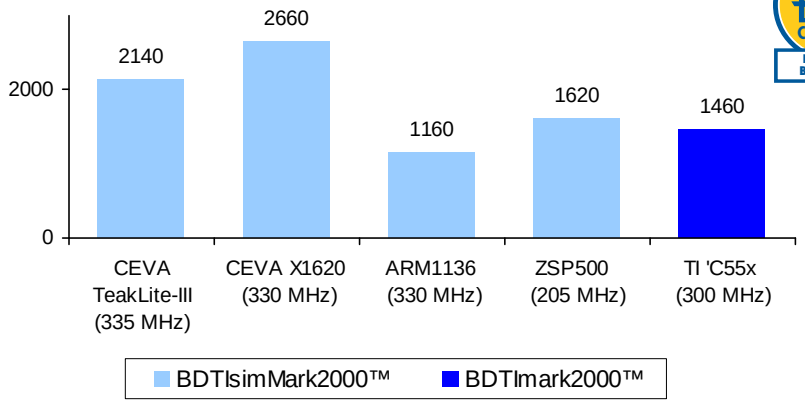

(The 'C55x is a chip; all others are cores)

based on BDTI's uniform conditions. In addition to this data, the vendor provides BDTI with a report detailing how this data was obtained. BDTI evaluates the vendor's results and methodology, and accepts the results only after confirming that they are in conformance with BDTI's specification.

Overview of Speed Results

As shown in Figure 2, the fastest processor in this group is the CEVA-X1620, followed by the TeakLite-III. The ARM1136, 'C55x, and ZSP500 are not nearly as fast as the TeakLite-III.

The score that a processor achieves on the BDTImark2000 and BDTIsimMark2000 metrics is proportional to both its clock speed and its architectural parallelism (i.e., how many operations it can execute in parallel). The ability to perform multiple MAC operations in parallel is often considered to be particularly important. Like the applications they represent, many of the BDTI DSP Kernel Benchmarks make heavy use of MAC operations.

It is important to note, though, that MAC throughput is not by itself a reliable predictor of performance. Many factors other than MAC throughput affect performance. This is particularly true for BDTI's single-sample benchmarks, such as the Single-Sample FIR filter benchmark, which tend to spend a minority of their cycles performing MAC operations, and for the Control and Viterbi benchmarks, which make almost no use of MACs at all.

This point is well illustrated by the benchmark results shown in Figure 2; all of the processors shown here are capable of executing two 16-bit MACs per cycle but their performance is quite different—even for those processors that are operating at nearly identical clock speeds. It is particularly interesting to note that TeakLite-III has nearly double the performance of the ARM1136 at roughly the same clock speed. As a CPU, the ARM1136 does not have many of the DSP-oriented features that TeakLite-III has, and its DSP performance reflects the lack. Of course, as a CPU the ARM11 has other advantages, such as the ability to run a full-featured operating system and the availability of a vast array of off-the-shelf software components for non-DSP tasks.

TeakLite-III is also significantly faster than the 'C55x, which is somewhat surprising because the 'C55x *is* a DSP processor; we'll discuss this result further in the next section.

Analysis of TeakLite-III Speed Results

TeakLite-III is able to provide strong DSP performance relative to the competitors shown here by using its specialized DSP-oriented features and instructions. The core's SIMD dual-MAC instructions are used fairly extensively, while specialized FFT and Viterbi instructions and accelerator hardware help it to achieve relatively low cycle counts on the FFT and Viterbi benchmarks. In addition, the TeakLite-III instruction set provides a "max2w" instruction that not only finds two maxima (for odd and even addresses), but also saves the locations of the maximas. As a result, the Vector Maximum benchmark cycle count is very low. Also, the ability to load unaligned 32-bit data helps the core perform well on BDTI's Bit Unpack benchmark.

As described earlier, TeakLite-III supports specific combinations of parallel instructions. Three of the parallel instructions are used in the FFT and Viterbi benchmarks, but most of the rest of them are not used in the benchmarks because their functionality is not well matched to the requirements of the benchmarks. Some of the parallel instructions (such as square and square accumulate) are designed for specific types of processing not used in our benchmark suite; others have strict rules

Figure 3. Energy Efficiency: BDTIsimMark2000™/mW (Higher is Better)

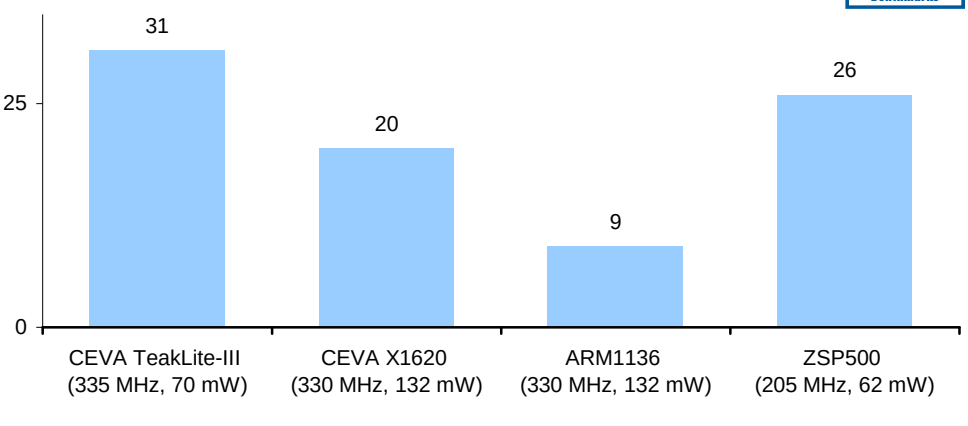

regarding their usage and are not flexible enough to be useful in our benchmarks. The IIR single-sample filter benchmark, however, does benefit from using parallel instructions, and achieves a very low cycle count.

The TeakLite-III inability to accumulate two 16×16-bit products separately increases the number of cycles required on the Real Block FIR, Complex Block FIR, and the LMS filter benchmarks. Furthermore, in some situations the TeakLite-III is unable to sustain two multications per cycle when performing 16-bit cross multiplies, which costs some cycles on the Complex Block FIR Filter.

Compared to the 'C55x, TeakLite-III has similar performance on the block FIR filters (as would be expected, since both are dual-MAC DSPs) but TeakLite-III has a significant advantage in the Vector Maximum, FFT, and Viterbi benchmarks because of its specialized instructions and hardware.

Both the superscalar ZSP500 and the VLIW-based CEVA-X1620 have somewhat more powerful instruction sets than TeakLite-III and have, overall, lower cycle counts across the benchmark suite. However, the ZSP500 runs at only 2/3 the clock speed of TeakLite-III, making it noticeably slower on our benchmarks.

Energy Efficiency

BDTI evaluates overall energy efficiency using the BDTImark2000/milliwatt metric. This is a composite performance metric based on a processor's typical energy use on the BDTI DSP Kernel Benchmarks. As with the speed results, power consumption is reported for a 0.13 µm process.

As shown in Figure 3, TeakLite-III has the best estimated energy efficiency of the cores compared here, with the ZSP500 following closely behind. This advantage should serve TeakLite-III well in mobile audio applications. In general, there is a trade-off between hardware specialization and energy efficiency; the more specialized a core is for its given workload, the better its energy efficiency will be (at the cost of flexibility and efficiency on other kinds of workloads). Based on these results, CEVA appears to have made a good trade-off between specialization (in terms of hardware and instructions) and flexibility for a range of DSP-oriented applications.

Area Efficiency

BDTI uses the BDTImark2000/mm^2 metric to evaluate a core's area efficiency. Figure 4 shows the area efficiency of the TeakLite-III core and selected competitors. As shown in this figure, TeakLite-III has the best area efficiency of the cores compared here. It is roughly half the size of its competitors, which, when combined with its relatively high signal processing speed, yields excellent area efficiency.

Memory Efficiency

The memory requirements of an application can have a significant impact on overall system cost. In addition, processors may experience significant performance degradation when application code and data do not fit in on-chip memory. Because of these and other factors, memory use efficiency is an important metric in processor selection.

In most applications, the control-oriented code consumes a much larger percentage of the total memory required than the signal-processing code. Therefore, when evaluating and comparing memory efficiency, BDTI assigns a much greater weight to Control benchmark results than it does to results for other benchmarks.

Based on our analysis of memory usage on the BDTI DSP Kernel Benchmarks, TeakLite-III will have memory efficiency that is very similar to that of the X1620, ZSP500, and ARM1136, and about 10% worse than that

Figure 4. Area Efficiency: BDTIsimMark2000/mm^2 (Higher is Better)

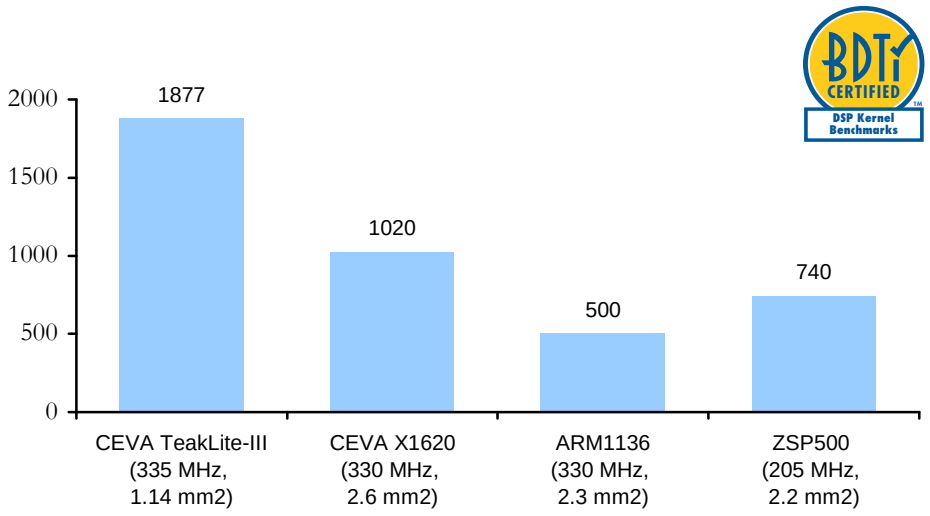

of the 'C55x. However, we should note that this evaluation is based on memory use in hand-coded BDTI Benchmark assembly routines, and therefore does not evaluate the efficiency of compiled code. We expect that most control-oriented software will be compiled rather than coded in assembly language; thus, the efficiency of the compiler will play a critical role in the memory efficiency of TeakLite-III (and other cores). BDTI did not evaluate the memory-use efficiency of the compiler.

Programming Model and Tools

This section provides brief comments about BDTI's experience of the programming effort required to develop optimized DSP software for the TeakLite-III and the available software development tools.

The TeakLite-III core is programmed in the manner typical for embedded processors, using a suite of development tools that include a C compiler, assembler, and instruction-set simulator, all accessed via an IDE. For this evaluation project, BDTI used the CEVA SmartNcode IDE version 9.1.7 and SmartNcode Debugger version 9.1.7.

Our overall experience with the IDE was quite positive. We found it easy to use, and the programmer has control over many aspects of the display—for example, the programmer controls how data is aligned and how the debugger windows are laid out, and the project configuration (i.e., window layout and compiler options) can be saved and used for other projects. The instruction-set simulator is described as cycle accurate, and in general it works reasonably well. However, the simulator sometimes did not report stall cycles correctly, which can make them more difficult to track down and eliminate—an important task when optimizing code for performance.

BDTI implemented and optimized the BDTI DSP Kernel Benchmarks almost entirely in assembly language and thus did not use the compiler much, though we did create a few simple programs in C. For these programs, we found that the compiler always generated correct code, but not necessarily well-optimized code. For example, the compiler is not able to generate the specialized FFT or Viterbi instructions.

For processors targeting audio applications, the availability of off-the-shelf optimized audio software—particularly codecs—is essential. CEVA provides a range of software modules optimized for the TeakLite-III-based CEVA-HD-Audio platform. These include:

- *DTS-HD: High Resolution Decoder, Master Audio Decoder, Express Decoder, DTS Core Decoder and Encoder, Neo6*
- *Dolby: Dolby Digital Decoder, Dolby Digital Plus Decoder, Dolby TrueHD Decoder, ProLogic-IIx Decoder, Dolby Digital Consumer Encoder*
- *LC-AAC/HE-AAC stereo/5.1 V1 and V2 Decoder and Encoder*
- *MPEG1/2/2.5 Layers 1/2/3 Decoder and Encoder*
- *WMA, WMA Pro Decoders*
- *RealAudio Decoder*
- *PCM Mixing, Sample Rate Conversion*

As of this writing, the proprietary codecs have not yet been certified.

These modules should be useful for system engineers implementing compression-based audio applications.

Closing Comments

Given CEVA's long history in licensable DSP cores (dating back to CEVA's origins as part of DSP Group in the early 1990's), it comes as no surprise that CEVA has fielded another quite competent DSP core. Indeed, CEVA's long track record is reassuring given that chip developers choosing a licensable core want to be able to count on their vendor being around to support them for the long term.

BDTI's benchmark results indicate that, compared to the other in-class processors benchmarked by BDTI, TeakLite-III provides strong DSP performance coupled with excellent area and energy efficiency. For applications with mid-range signal processing requirements and stringent die size or energy consumption constraints, TeakLite-III deserves serious consideration.

The TeakLite-III software development tools also appear to be quite competent. (However, as mentioned earlier, BDTI did not have the opportunity to evaluate the TeakLite-III compiler in depth. Compiler efficiency is, increasingly, a key requirement for embedded processors, and system designers would do well to evaluate the associated compiler before choosing a core.)

Although general-purpose CPU cores are increasingly capable of handling DSP tasks, TeakLite-III shows that for demanding, DSP-centric applications, more specialized architectures can deliver significant performance and efficiency advantages.