

Optimizing the Implementation of Dolby Digital Plus in SoC Designs

Chris Cavigioli and Brett Miller, MIPS Technologies
Roger Dressler and Rob Hislop, Dolby Laboratories

January 2007

MIPS Technologies and Dolby Laboratories have collaborated to provide a rapid and low-risk methodology for deploying Dolby Digital Plus technology to the set-top box and HD DVD/Blu-ray Disc markets.

Introduction

Dolby[®] Digital Plus, the newest generation of Dolby Digital, the standard-setting audio technology from Dolby Laboratories, is making its way into next-generation home entertainment applications, bringing superior sound quality, more efficient audio compression, and an improved user experience to the home.

Dolby Digital Plus, also known as Enhanced AC-3 (E-AC-3), was developed to address the changing requirements of two burgeoning consumer markets. For the emerging HD DVD and Blu-ray Disc optical disc players, the data compression capabilities of Dolby Digital Plus allow movie studios to combine a superior audio experience with high-definition video. The technology also enhances the latest digital TV set-top boxes (STBs), enabling broadcasters to deploy offerings with lower bit rates, reducing costs and adding flexibility and value for consumers.

While there are tremendous advantages to using Dolby Digital Plus, there are also design considerations that system-on-a-chip (SoC) teams need to understand. The rigorous approval testing that is required at Dolby Laboratories before a chip or a system can use the Dolby Digital Plus logo is substantial. Working with Dolby Laboratories, MIPS Technologies has developed an optimized, tested version of Dolby Digital Plus that will run on any of its 32-bit synthesizable processor cores—significantly slashing audio subsystem development time. This approach also lowers the overall cost of the SoC by eliminating the need for an expensive digital signal processor (DSP).

This white paper outlines the benefits of Dolby Digital Plus for select key consumer audio markets and details the processes behind the collaborative effort between MIPS Technologies and Dolby Laboratories.

Emerging Standards for the Home Entertainment Market

HD DVD and Blu-ray Disc

Dolby Digital Plus is mandatory in the HD DVD specification and is defined as optional in Blu-ray Disc. The specifications also allow for multiple independent streams so a movie can be shipped with multiple independent soundtracks or support multiple languages.

For the first time, HD DVD and Blu-ray Disc players allow external streams from the Internet to be mixed with the native audio on the disc after the disc has been shipped to the consumer. For example, a high-definition video stream can be accompanied by an audio track mastered on the disc that is simultaneously mixed with a more recent director's commentary streamed from a studio website.

Next-Generation Set-Top Boxes

Dolby Digital Plus is also vital to the next generation of set-top boxes. High-definition TV depends on high-quality sound for a complete experience. Broadcasters and network operators are all looking at ways to dramatically improve their offerings.

The greater compression of Dolby Digital Plus allows 5.1 surround sound using as few as 224 kbps, a significant reduction from the current 384–448 kbps data rate typically used for Dolby Digital (AC-3). This provides an opportunity for broadcasters to enhance their services by allowing value-added services for the hearing or visually impaired. Broadcasters using MPEG stereo at 160–256 kbps can upgrade to 5.1-channel surround sound using essentially the same bandwidth (see Figure 1).




Figure 1: The Advantages of Upgrading to Dolby Digital Plus for Broadcasters

Dolby Digital Plus maintains backward compatibility with the worldwide installed base of more than 36 million Dolby Digital audio/video receivers using a S/PDIF connection (see Figure 2a).




Figure 2a: The S/PDIF Connection from STB or Optical Disc Player to an A/V Receiver

Fully decoding the stream to linear PCM and then re-encoding it to Dolby Digital would require significant processing power and would reduce the audio quality. Using an efficient low-latency, low-complexity conversion process, Dolby Digital Plus streams are converted into Dolby Digital streams more easily and with less impact on audio quality.

With the high-speed HDMI™ interface becoming increasingly common on newer STBs, the Dolby Digital Plus streams (either coded or uncoded) can be transmitted directly to a new generation of A/V receivers, resulting in audio of a higher quality than ever before (see Figure 2b).




Figure 2b: The HDMI Connection on a Set-Top Box Allows a Full Dolby Digital Plus Stream to Be Transmitted Directly

Dolby Digital Plus Software Development

Some system developers use a dedicated audio DSP and adapt the audio code from the floating-point reference code that is supplied by Dolby Laboratories. This approach can be arduous and expensive considering the cost of the extra silicon, DSP royalties, and board tests. Working with a different set of development tools for the DSP design in addition to the system design also increases the time it takes to develop the implementation. The end unit then has to be extensively tested and approved by Dolby Laboratories to attain the Dolby Digital Plus logo.

A lower-cost, more flexible approach is to use a general-purpose processor core for the algorithms. The audio processing can be subsumed by the host CPU or

moved to a second processor in a system chip, providing the Dolby Digital Plus capability at a nominal cost. Fewer components are needed; less board space and less power are used. Because it is general-purpose, the core can be used for other tasks in addition to decoding or converting the audio stream.

This approach requires a high-performance processor core and code that has been optimized specifically for that core and approved by Dolby Laboratories. MIPS Technologies has worked closely with Dolby Laboratories on the implementation of the Dolby Digital Plus decoder. For designers or users of MIPS-Based™ SoCs, this significantly reduces the risk, timescale, and cost of adding Dolby Digital Plus to the end product.

Implementation

The Dolby Digital Plus reference software is written in floating-point C code and delivered to the licensee for implementation in a device or a DSP that is then sent to Dolby Laboratories for approval testing.

Dolby Laboratories maintains a high standard of quality, so the testing procedure is correspondingly extensive.

To help ensure that the approval process goes smoothly, Dolby Laboratories provides two sets of tests to the implementer for use during development: a set of test vectors to rigorously exercise each part of the algorithm at a low level and quickly compare the output of a decoder under test (DUT) with the output of a reference implementation, and a more rigorous set of tests to verify the final implementation using Audio Precision® test equipment. The Audio Precision tests essentially duplicate the testing Dolby Laboratories will perform. Passing both sets of tests in-house is required prior to submission to reduce risk and makes it more likely that the implementation will pass at Dolby Laboratories.

As noted, MIPS Technologies has developed an optimized version of the code to streamline the process and to provide a high-quality implementation. This code, which has been approved by Dolby Laboratories, is provided to developers who have licensed both the Dolby Digital Plus technology and MIPS processor technology. While a MIPS-Based SoC implementation still must go back to Dolby Laboratories for approval, using the pre-approved code significantly reduces risk for the implementer and helps ensure a quicker turnaround for approval at Dolby Laboratories (see Figure 3).

Dolby Laboratories benefits in several ways, as most of the STB and optical disc SoCs are based on industry-standard MIPS processor cores.




Figure 3: Approval Process for Dolby Digital Plus System Chips and Equipment

Using the MIPS-optimized version of the code can shave months from the development time of a typical project and accelerate the approval cycle. In a similar fashion, MIPS Technologies also offers optimized and approval-tested code for Dolby Digital (AC-3), Dolby Pro Logic® II, and other popular audio algorithms such as MPEG-2, AAC, MP3, and SRS TruSurround XT®.

Optimizing for MIPS32 Implementation

In embedded systems, the most efficient implementation of an audio algorithm uses fixed-point arithmetic rather than floating-point arithmetic. Therefore, one of the first steps in optimizing the Dolby reference implementation is to convert the floating-point reference code.

A high-quality conversion requires choosing a fixed-point representation (or potentially multiple representations at different processing stages) that finds an optimal balance of dynamic range and precision. This allows the algorithm to handle any potential input without exceeding the dynamic range of its representation while still providing sufficient precision. Just as important, choosing the optimal representation avoids the need for saturation checks after each operation, significantly improving run-time performance.

MIPS Technologies accomplishes this conversion process using an extensive simulation and testing procedure, ensuring that the implementation will meet the stringent audio-quality standards established by Dolby Laboratories. An abstract C++ fixed-point class tracks the required dynamic range and precision of data throughout the algorithmic processing. At the end of the simulation process, fixed-point representations and operations are chosen and hardwired back into a pure C implementation.

When choosing fixed-point representations, MIPS takes full advantage of its architecture. By carefully choosing representations, the 32-bit general-purpose registers (GPRs) of the MIPS32® architecture allow for a similar capability as a DSP accumulator with guard bits, even with GPR-to-GPR operations. In some cases, where 32 bits is still insufficient, the HI/LO register pair of the accumulator can be used as a 64-bit accumulator to provide extra precision or dynamic range as needed.

Converting the source code to fixed point is already one form of optimization. And given the fixed-point C source code, the MIPS Software Development Environment (SDE) compiler tool chain is capable of producing highly efficient code, scheduled for a specific MIPS core pipeline without additional effort. However, some of the fixed-point operations and critical DSP routines, such as IMDCT (inverse modified discrete cosine transform) and window overlap/add, can further benefit from architecture-specific optimizations that the compiler alone is not capable of generating. Such optimizations require a detailed understanding of both the MIPS architecture and the structure of the DSP routines employed in Dolby Digital Plus.

Soft Audio Interface Integration

After producing a fully optimized fixed-point implementation of the Dolby Digital Plus code, MIPS encapsulates the decoder-converter with its Soft Audio Interface (SAI), a standard API supported by all MIPS-optimized audio codecs. The SAI abstracts many of the complexities of particular algorithms and allows developers to think of them in terms of components that can be readily added or removed from an application. By standardizing the way in which the MIPS Soft Audio algorithms are used, the SAI also facilitates the reuse of application-level code for different or changing audio processing tasks, and makes it easy to cascade multiple audio codecs using shared intermediate memory buffers (see Figure 4).




Figure 4: Converting Floating-Point Dolby Code to Highly Optimized Fixed-Point MIPS32 Code

All this allows the system designer significant flexibility. The codec can be run on the main system CPU or on a second core used as an audio coprocessor. The size of the core and its supporting memory is an issue for the system architect.

Using a general-purpose MIPS core means it can be used for a wide range of other functions as well as Dolby Digital Plus and other audio codecs.

This approach also allows the system design to be scalable. A small low-cost, entry-level 4KETM core can be used for a low-bandwidth implementation of Dolby Digital Plus, while the latest multi-threaded 34KTM core could be used for a high-definition implementation and a host controller (see Figure 5). The multi-threading allows the 34K core to behave as two virtual processors within the single core by optimizing the use of the instruction pipeline and the memory to fully maximize the hardware. The two processors look independent to the main system, so one can handle the host operating system while the other handles the audio and video subsystems.

Finally, there is a distinct advantage to the developer designing a multi-core system (a host processor and audio coprocessor, for example) to utilize a common development environment. Using the same set of development tools makes it easier for teams to work together, to debug the system, and also simplifies load balancing.




Figure 5: The Tested, Optimized Dolby Digital Plus Code Can Be Used Across the Entire Array of MIPS Cores

Conclusion

Dolby Digital Plus is a new audio coding technology that represents significant value to broadcasters in the STB market and is specified as a standard audio format for the high-definition video disc markets. Dolby Digital Plus saves 33 percent bandwidth for broadcasters, giving them lower transmission costs or the ability to deploy value-added services on the unused bandwidth. Broadcasters using MPEG stereo can now transmit 5.1-channel surround sound in virtually the same bandwidth using Dolby Digital Plus. In the blue-laser optical disc market, Dolby Digital Plus is mandatory for HD DVD and optional for Blu-ray Disc.



MIPS is the number-one architecture in the STB, DTV, and DVD recorder markets today. In order to facilitate the rapid deployment of Dolby Digital Plus across the variety of MIPS-Based SoCs in those markets, MIPS Technologies has licensed the Dolby Digital Plus technology and already completed the tedious task of optimizing it in fixed point for the whole family of MIPS processors. Now, SoC designers have a shortcut to get to market with less risk. They get the reference code from Dolby Laboratories, but they can confidently use the code from MIPS Technologies that has already been optimized for the MIPS architecture and has already been tested and approved by Dolby Laboratories. Final SoCs are more likely to pass Dolby Laboratories approval testing quickly.

With the cooperation between MIPS Technologies and Dolby Laboratories, the industry benefits by saving development costs and reducing risk, bringing Dolby Digital Plus, and its benefits, to the retail consumer faster.