1 Introduction

The eco32 processor family was defined with digital audio as a major target application on top of general purpose embedded control and computing. This mp3 decoder test was implemented to fine tune the eco32 ISAs for excellent performance and computational precision in digital audio. Mp3 decode was chosen because it is the most wide spread digital audio applications and it requires a fair mix of control code and DSP code. To achieve superior results a processor must be good in both categories. Another challenge is to achieve high computational precision at the same time with high performance using a 32-bit fixed point implementation.
The mp3 decoder is entirely written in assembly. A C compiler for the eco32 processor was not yet available at the time the decoder was implemented. The performance critical subroutines are highly optimized and follow the ABI (Application Binary Interface) that is planned for eco32 C compiler implementations. For commercial mp3 decoders these optimized routines can be reused.

2 Decoder details

2.1 Features

The decoder implements most features that are important for a main stream commercial mp3 player product:

- Decodes single-channel and dual-channels MPEG layer III audio streams.
- Supports 32KHz, 44,1kHz and 48kHz sample rates
- Supports all layer III fixed bit-rates from 32 kbits/s to 320 kbits/s
- Supports VBR (Variable Bit-Rates)
- Supports joint stereo streams, middle/side and/or intensity
- Outputs 16-bit rounded and clipped samples
- Fully compliant decoder according to the ISO/IEC 11172-4 precision test

The following features are not supported at this time because they were not relevant for the goals of this test bench:

- No support for layer I and layer II streams
- No support for free format bit-rates
- No support for MPEG2 and MPEG2.5 sample-rates
- CRC check not implemented yet (but streams with CRC are decoded)

The decoder also implements some error robustness and recovery features. They were necessary to be able to decode some of the streams used for testing. Because mp3 is not very precisely defined and standardized many streams would fail or crash the decoder without these error detection and recovery mechanisms. The following mechanisms are in place:

- Illegal header parameters -> stop decoder
- Illegal Huffman table -> discard granule and mute output
- Bits left in the bit reservoir after Huffman decode -> read remaining bits from stream and discard
- Under-run of the bit reservoir during Huffman decoding: -> return extra bits to stream
- Main data value too large (points beyond buffered stream) -> discard frame and wait for legal value

2.2 Precision

The ISO/IEC 11172-3 standard defines a compliance test for the precision of layer III decoders. The output of the decoder for a single channel sine sweep from 20Hz to 10kHz is compared against a reference file with full precision 24-bit samples. Three levels of compliance are defined: 1. fully compliant, 2. limited accuracy and 3. not compliant. Two parameters are relevant for the test:

- The maximum absolute difference between the decoder output samples and the reference samples
- The RMS (root mean square) value of the absolute difference signal between the decoder output and the reference samples, calculated over the entire test stream

The sine sweep test stream has 216 frames and produces 248.832 samples. The output of the eco32 mp3 decoder after clipping and rounding to 16-bit samples was used for the compliance test. For a “fully compliant” result, the maximum absolute difference between any sample of the decoder and the corresponding reference samples must be $< 2^{-14}$ ($\sim 6.104*10^{-5}$) and the RMS over the entire difference signal of the clip must be $< 2^{-15}/\sqrt{12}$ ($\sim 8.810*10^{-6}$). The test result for the eco32 mp3 decoder is $1.752*10^{-5}$ for the maximum difference and $8.614*10^{-6}$ for the RMS over the difference signal. With that the eco32 mp3 implementation is a fully compliant decoder.

2.3 Performance

The differentiating feature of the eco32 architecture with respect to digital audio is the excellent performance of both ISAs combined with high computational precision. For bith base and DSP ISAs this is exceptional because they look like normal 32-bit RISC processors and not like dedicated audio DSPs.

To enable fair comparison with competing solutions the environment for performance measurements must be described especially any acceleration hardware but also any software functions that count into the measured numbers and are not necessarily required for a pure decoder software.

A single hardware assist feature is used that must be mentioned: the eco32 mp3 decoder implementation uses a “bit string reader” peripheral that allows access to non aligned bit strings from 0 to 24-bit length. Main reason for using this “accelerator” is the reuse of a video decoder test bench that was developed before the mp3 decoder. Also, in a combined audio/video decoder application the peripheral can be used for both audio and video decoding. However the acceleration effect is much more significant for video where the symbol decode rates are much higher and the associated bit-string reads represent a significant share of a software decoder’s performance requirements. For an mp3 audio decoder the acceleration effect is minor. Exact measurements have not been done, but a fair assumption is that the effect is well below 5% and therefore not significant for the performance results.

On the other hand the performance measurements include two software functions that are not necessarily part of an mp3 decoder, but could as well be part of a player application:

- A 2kBytes buffer in the data memory space is used as intermediate buffer between reads from the stream file and bit string accesses of the decoder. The maintenance of this buffer (fullness checking, refilling through a software loop) is included in the performance numbers.
- Output samples are rounded and clipped to 16-bit values before they are written to the PCM output buffer. This computation step is included in the performance numbers

Performance has been measured on a number of audio streams from various mp3 encoders. The streams have different bit-rates and stereo modes. The tests have been done in an environment with zero wait state instruction and data memories.

The table below shows the performance measurement results. For each stream the performance of ISS simulations of the two ISAs and of the eco32bl and eco32dl hardware implementations are given. Performance is expressed as average MHz numbers with two digits after the decimal point. The following formula is used to calculate the performance numbers:

Performance[MHz] = ((cycles/frames) * (samplerate/1152)) / 1000000

For most of the test clips the simulations were run for 1000 frames. For clips with < 1000 frames lengths the simulations were run for the entire clip.

stream	bit rate [kBps]	properties	eco32b	eco32bl	eco32d	eco32dl
1	128	MS stereo	9.87	19.40	7.47	7.94
2	112	MS stereo	9.46	18.79	7.16	7.58
3	128	MS stereo	9.77	19.16	7.41	7.88
4	VBR (190 aver.)	MS stereo	10.42	20.18	8.01	8.63
5	VBR (142 aver.)	MS stereo	10.05	19.49	7.65	8.17
6	128	dual channel	9.77	19.04	7.42	7.95
7	96	MS / int. stereo	8.63	17.01	6.48	6.87
8	160	MS stereo	10.61	20.50	8.12	8.70
9	192	stereo	10.24	19.82	7.90	8.57

The eco32bl is the low end member of the eco32 32-bit processor family. It only has a single 16*16-bit hardware multiplier, 32*32-bit multiply and multiply/accumulate instructions take multiple (up to six) clock cycles. As a result the mp3 decoder performance is significantly lower (about factor 2) than the theoretical performance of the eco32b ISS model. But, considering that the eco32bl is an entry level 32-bit general purpose processor with only about 35kgates size the mp3 performance is still very good. For audio centric applications that require either high performance or very low power consumption it is recommended to use the eco32dl.

The eco32dl has a 32*32-bit hardware multiplier and multiply as well as multiply/accumulate instructions are executed in one cycle effective. In addition the DSP ISA extension supports memory source operands. With these features the eco32dl achieves excellent mp3 decoding performance and can easily compete with dedicated audio DSPs. With only about 60kgates size the eco32dl is mainly targeted at low power and low cost applications.

2.4 Power 4/3 calculation

In mp3 decoding, the formula for the de-quantization of decoded Huffman symbols requires the calculation of the power**4/3 value of the decoded symbols. The range of input values is from 0 to 8206 (integers). Mp3 decoders for PCs typically use a table lookup to determine the symbol**4/3 value. With 32-bit output values this requires a table of about 32kBytes which is no problem for a PC application. But for embedded applications, especially low-cost, portable mp3 players a 32kBytes table is a problem for a software application that otherwise requires only about 8kbytes of total constant memory space. Exact computation of the symbol**4/3 value would be very compute intensive and is therefore unpractical.

The eco32 mp3 decoder uses a compromise between performance and memory space requirement. For input values from 0 to 63 a lookup table with precise pre-computed values is used. For the statistically rare input values > 63 piecewise approximation using 3rd order polynomials is used. This method requires a total table space of 556 bytes and the approximation method is good enough to achieve full compliance in the ISO/IEC 11172-4 precision test.

2.5 Memory footprint

The table below shows the total memory space requirements of the mp3 decoder software. The “text” (instruction code) numbers include debugging routines.

ISA	Text (instruction code)	Variable data	Constant data
eco32b	9188 Bytes	~24 kBytes Incl. stack	8456 Bytes
eco32d	9056 Bytes		