Power Consumption of Customized Numerical Representations for Audio Signal Processing

**Roger Chamberlain, Yen Hsiang Chew,
Varuna DeAlwis, Eric Hemmeter,
John Lockwood, Robert Morley,
Ed Richter, Jason White, Huakai Zhang**

Washington University
Campus Box 1115
One Brookings Dr.
St. Louis, MO 63130-4899

Power Consumption of Customized Numerical Representations
for Audio Signal Processing [*]

Roger Chamberlain, Yen Hsiang Chew, Varuna DeAlwis, Eric Hemmeter,
John Lockwood, Robert Morley, Ed Richter, Jason White, and Huakai Zhang
School of Engineering and Applied Science
Washington University, St. Louis, Missouri

Abstract

One of the major technical issues facing the designers of modern, hand-held, portable, digital systems is the need to minimize the power consumption of the system to prolong battery life. Many of these systems perform signal processing functions on audio signals (e.g., communications systems, hearing aids, MP3 players, etc). As new signal processing techniques are proposed, the computational requirements invariably grow, putting additional pressure on power consumption. In this work, we investigate the use of non-standard numerical representations for processing of audio signals, showing how the power consumption can be lowered for audio signal processing while maintaining (and even improving) overall signal quality.

Standard numerical representations used for signal processing applications include fixed-point representations (typically 16 bits) and floating-point representations (either 32- or 64-bit IEEE standard). The choices of representation available to system designers are much more based upon historical convention than the specific needs of the application. Our purpose is to investigate the implications of deviating from these standard representations and designing a system with a numerical representation tailored to the specific application.

Focusing our attention on audio signals that communicate human speech, a dynamic range of approximately 100 dB and a signal-to-quantization-noise ratio (SNR) of approximately 30-35 dB have been shown to be adequate [1]. As a representative computation, we investigate the power consumption associated with a series of multiply-accumulate operations. The multiply-accumulate (MAC) is the most common computation in audio signal processing.

For standard, fixed-point, linear, numerical representations executing MAC operations, the multiplication is significantly more expensive (in terms of power) than the accumulation (by about an order of magnitude). This has motivated previous investigations into logarithmic representations [2], exploiting the fact that the multiplication operations can be implemented using an adder:

$$\log (a*b) = \log (a) + \log (b).$$

This savings must be traded off against a larger power consumption in the accumulation operations, which have been implemented as look-up tables (LUTs) in previous work.

If a logarithmic representation gives a significant power savings in the multiplication, but costs power for the accumulation, a numerical representation that is partially logarithmic and partially linear has the potential to achieve balanced power consumption across both operations. This is precisely the description of a floating-point representation, where the exponent represents the logarithmic portion and the mantissa represents the linear portion. Rather than IEEE standard floating-point representations, however, we are interested in representations designed to more closely match the requirements of the audio signal processing application.

In this study, we compare the power consumption of a 16-bit linear representation with several different floating-point representations (4- to 6-bit exponent and 4- to 6-bit mantissa) and a 9-bit logarithmic notation. Each representation is tailored to meet the dynamic range and SNR requirements described

---

above. For each representation, we design a hardware MAC unit in the VHDL language and perform a standard-cell synthesis, layout, and place-and-route targeting the AMI Semiconductor 0.5 micron VLSI integrated circuit process. The standard cell library is made available through the Mentor Graphics Higher Education Program. The resulting design is simulated using the Mentor Graphics MACH-PA power analysis tool, with random input vectors. The simulation output both verifies correct operation of the circuit and provides information on power consumption.

Figure 1 plots the SNR over the range of valid numbers that can be represented for three different floating-point representations. Here, the "s-e-m" notation refers to "s" sign bits, "e" exponent bits, and "m" mantissa bits. For example, "1-6-4" is a representation with 1 sign bit, 6 bits of exponent, and 4 bits of mantissa. Notice that although each representation has a total of 11 bits, each representation has different SNR and dynamic range properties. For the 1-6-4 representation, the smaller number of bits in the mantissa implies a lower SNR, which ranges between 20 and 25 dB, while the larger number of bits in the exponent implies a larger dynamic range (385 dB in this case). For the 1-4-6 representation, the SNR ranges from approximately 32 to 37 dB, however, the dynamic range is considerably smaller at 96 dB.
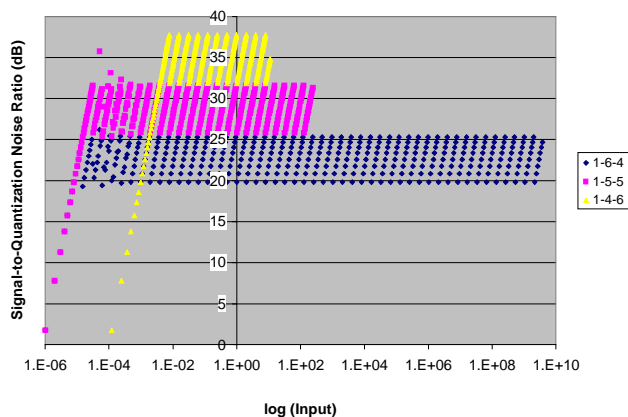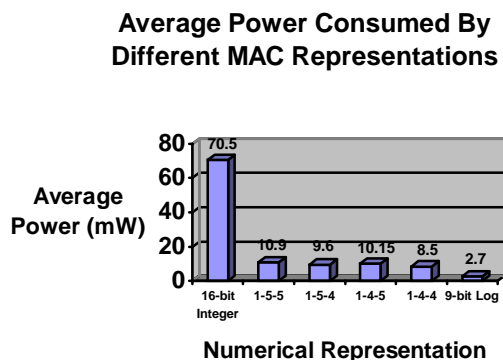

Figure 1. Numerical properties.

Figure 2. Power consumption.

Figure 2 shows the average power consumption of the MAC unit when executing a series of operations. The results show a significant power savings (approximately an order of magnitude) using both the floating-point representations and the logarithmic representation. This is primarily due to the ability to either eliminate (in the case of the logarithmic representation) or significantly reduce the size of the hardware multiplier required as part of the MAC unit. We will present both novel techniques for implementing the accumulation function with a logarithmic representation as well as the power consumption associated with a wider set of numeric representations.

Although the specific results presented here are limited to a small range of numerical representations, the general message is much broader. With the advent of reconfigurable logic, rigid, fixed-function computational hardware is no longer a necessity of modern digital system design, and significant power savings can result if the requirements of the application are used to specify the properties of the numerical representation. While this work was initially motivated by the needs of digital hearing aids [3,4], it is applicable to a wide variety of digital signal processing applications.

[1] S.D. Soli, "Offline simulations of logarithmic FIR filters: Measurements of processing distortion and perceptual tests of its audibility," Technical report, Bioscience Laboratory, 3M Center, October 1988.
[2] T.J. Sullivan, "Estimating the power consumption of custom CMOS digital signal processing integrated circuits for uniform and logarithmic number systems," D.Sc. dissertation, Washington University, 1993.
[3] J.L. Goldstein, M. Valente, and R.D. Chamberlain, "Acoustic and psychoacoustic benefits of adaptive compression thresholds in hearing aid amplifiers that mimic cochlear function," *J. Acoust. Soc. Am.* **109**:2355(A), 2001.
[4] R. Chamberlain et al., "Novel Numerical Representations for Low-Power Audio Signal Processing," to appear at the *Int'l Hearing Aid Research Conf.*, August 2002.