The "ARMAdillo" Coefficient Encoding Scheme for Digital Audio Filters

Dave Rossum, E-mu Systems
1600 Green Hill Road, PO Box 660015
Scotts Valley, California, 95066-0015 (408) 438-1921

In the design of VLSI circuits to implement digital filters for electronic music purposes, we have found it useful to encode the filter coefficients. Such encoding offers three advantages. First, the encoding can be made to correspond more properly to the "natural" perceptual units of audio. While these are most accurately the "bark" for frequency and the "sone" for loudness, a good working approximation is decibels and musical octaves respectively. Secondly, our encoding scheme allows for partial decoupling of the pole radius and angle, providing superior interpolation characteristics when the coefficients are dynamically swept. Thirdly, and perhaps most importantly, appropriate encoding of the coefficients can save substantial amounts of on-chip memory. While audio filter coefficients typically require twenty or more bits, we have found adequate coverage at as few as eight bits, allowing for a much more cost effective custom hardware implementation when many coefficients are required. We have named the resulting patented encoding scheme "ARMAdillo."

Our implementation of digital audio filters is based on the canonical second order section whose transfer function should be familiar to all:

$$H(Z) = a_0 \frac{1+\frac{a_1}{a_0}z^{-1}+\frac{a_2}{a_0}z^{-2}}{1+b_1z^{-1}+b_2z^{-2}} \quad [1]$$

While dealing with poles and feedback (b_n) coefficients, the comments herein apply as well to zeroes and feedforward coefficients (a_n/a_0) when the gain (a_0) is separated as shown above.

Noting that the height of a resonant peak in the magnitude response produced by a pole is approximately inversely proportional to the distance from the pole to the unit circle, we can relate the height ρ of this resonant peak in dB to the pole radius R:

$$\rho \approx 20 \log_{10} \frac{1}{1-R} \quad [2]$$

To accomplish our purpose of a coefficient encoding which will be linear in ρ we encode b_2 in a variable k_2 which varies from 0 to 11:

$$b_2 \approx 1 - e^{-k_2}$$

Noting that $b_2 = R^2$ for complex poles, and relying on the fact that for significant resonance both R and b_2 will be near unity, we can derive:

$$\rho \approx 20 \log_{10} \frac{1}{1-\sqrt{1-e^{-k_2}}}$$

but note that for small ε, $\sqrt{1-\varepsilon} \approx 1 - \frac{\varepsilon}{2}$, so:

$$\rho \approx 20 \log_{10} 2e^{k_2}$$

Simplifying:

$$\rho \approx 8.68\, k_2 + 6.02$$

thus demonstrating k_2 will indeed approximate a decibel encoding of the filter resonance to heights of about 100 dB with the chosen range. Note that the approximations break down for k_2 near 0, but that its span guarantees access to the entire stable span of b_2 and R.

In practice, we use an octave shift encoding scheme to reduce the exponential decoding to a few logic gates and a barrel shifter. The errors introduced by such a scheme are in the neighborhood of ±1 dB.

We similarly encode coefficient b_1 to linearly represent musical octaves of frequency. Because, for complex poles, b_1 depends on both the radius R and the angle θ, our encoding scheme uses our previously defined variable k_2 to reduce this mutual dependency as well as further optimizing the encoding efficiency. We encode b_1 using a new variable k_1 which varies from 0 to 11:

$$b_1 \approx -2 + e^{-k_2} + 4e^{-k_1}$$

To show that k_1 is indeed approximately linear in musical octaves for the filters of interest, we must define the "musical octave number" Ω varying from zero to ten which is logarithmic in resonant frequency based on the pole angle θ and sample rate F_s:

$$\Omega \approx \log_2 \frac{\theta F_s}{40 \pi}$$

Because the majority of the audio band is in the lower octaves, we can use approximations which rely on small θ. From the fact that for complex poles $b_1 = -2 R \cos\theta$, we can derive:

$$\Omega \approx \log_2 \frac{F_s \operatorname{acos} \frac{-b_1}{2\sqrt{b_2}}}{40 \pi}$$

Substituting our encoded variables k_1 and k_2:

$$\Omega \approx \log_2 \frac{F_s \operatorname{acos} \dfrac{1 - \frac{e^{-k_2}}{2} - 2e^{-k_1}}{\sqrt{1 - e^{-k_2}}}}{40\pi}$$

Since for small ε, $\sqrt{1-\varepsilon} \approx 1 - \frac{\varepsilon}{2}$ and $\frac{1}{1-\varepsilon} \approx 1+\varepsilon$:

$$\Omega \approx \log_2 \frac{F_s \operatorname{acos}((1 - \frac{e^{-k_2}}{2} - 2e^{-k_1})(1 + \frac{e^{-k_2}}{2}))}{40\pi}$$

Applying that for small ε, $\varepsilon^2 \approx 0$:

$$\Omega \approx \log_2 \frac{F_s \operatorname{acos}(1 - 2e^{-k_1})}{40\pi}$$

Noting that for small ε, $\operatorname{acos}(1-e) \approx \sqrt{2e}$:

$$\Omega \approx \log_2 \frac{F_s \sqrt{4e^{-k_1}}}{40\pi}$$

Manipulating the logarithms, we find (for F_s=40kHz):

$$\Omega \approx \frac{-k_1}{2\ln 2} + \log_2 \frac{F_s}{20\pi} \approx \frac{-k_1}{1.38} + 9.31$$

We have thus shown k_1 is seen to be linearly related to frequency expressed in musical octaves.

The final proof of the encoding scheme lies in the effectiveness of the implementation used with "real" filters. We have found a useful graphical analysis tool to be the plot of the poles translating the radii from R to R' and the angles θ to θ' (excluding all θ below $\pi/1024 \approx 20$ Hz) such that:

$$R' = 20\log_{10}\frac{1}{1-R}$$

and

$$\theta' = \frac{\pi(10+\log_2\frac{\theta}{\pi})}{10}$$

Such a plot maps the poles onto a nominally log/log semicircle, and even coefficient density indicates an even perceptual mapping.

Figure 1 shows such a plot of unencoded coefficients, and Figure 2 shows a plot at the same coefficient quantization using ARMAdillo encoded coefficients. For interest, Figure 3 shows the encoded coefficients plotted in the traditional polar plot, showing the high coefficient density at the low frequencies and high resonances, as well as explaining our peculiar name for the scheme.

The physical implementation in a VLSI circuit confirms the success of the method. It linearly interpolates the coefficients in the encoded space at the sample rate; the resulting dynamic filters are smoothly controllable over a wide range with a pleasing perceptual mapping.

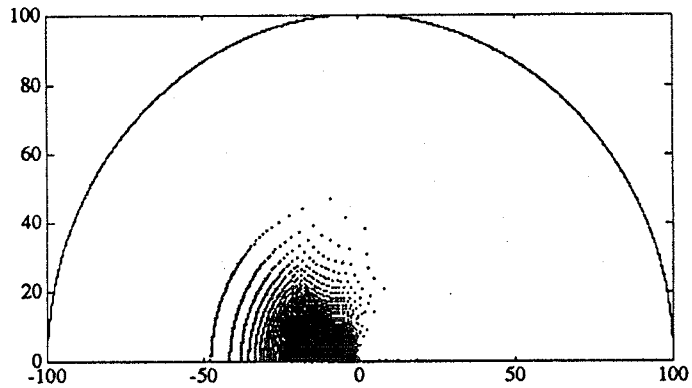

Figure 1 - Log Plot of Unencoded Coefficients

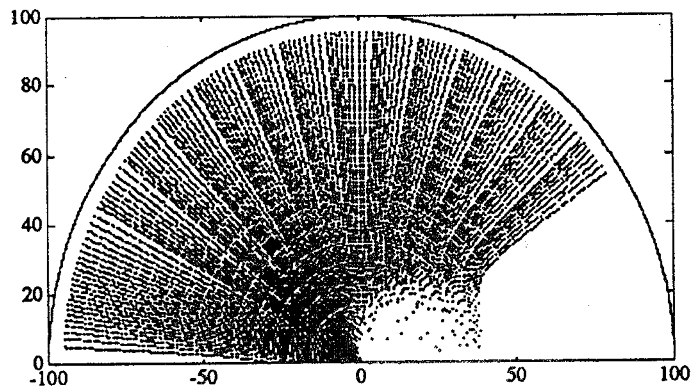

Figure 2 - Log Plot of Encoded Coefficients

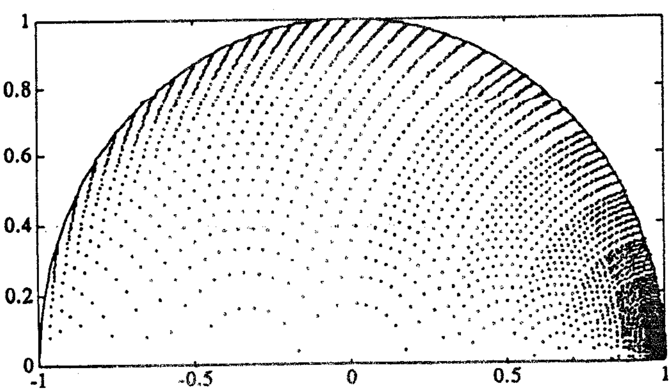

Figure 3 - The "ARMAdillo" Plot