

WHITHER DITHER: Experience with High-Order Dithering Algorithms in the Studio

By:
James A. Moorer
Julia C. Wen

Sonic Solutions
San Rafael, CA
USA

An ever-increasing number of recordings are being made at resolutions greater than 16 bits and more than 100 dB of dynamic range. Since the ultimate product of most of these recordings is a Compact Disc, there is some question of how to reduce the resolution of the material down to the required 16 bits. Noise shaping and dithering have been proposed as methods of doing this. These methods are attractive since they place the irreducible quantization error in a portion of the spectrum where the ear is relatively insensitive. We show that the success of this process depends greatly on the ultimate use and listening environment for the material. This moves the choice of what algorithm to use more into the domain of the mastering engineer than of the scientist. We conclude that it is thus not possible to "psycho acoustically optimize" an algorithm, contrary to the claims of some manufacturers and researchers. We recommend that equipment manufacturers present the client with a wide variety of algorithms so that the process can be controlled based on what end result is desired.

BACKGROUND:

Truncation of a high-resolution digital audio signal down to, say, 16 bits produces an error that cannot be reduced in magnitude. Moreover, the nature of the error is such that it is perceived as having rather annoying quality. The error is highly modulated by the signal, it is rich in overtones, and it masks the desired signal. There are a number of different ways the reduction to 16 bits may be done. We will discuss here forms that follow the processing shown in Figure 1. We shall take the box marked DITHER to be a generator of random noise with either a uniform distribution function of one least significant bit in amplitude or a triangular distribution function of amplitude twice the least significant bit. We will identify these as "UPDF" or as "TPDF." There are four principal variations on this scheme that we will identify as follows:

- Q1: Pure truncation** - This is identified by $DITHER = 0$ and $F(z) = 0$.
That is, it is pure quantization. It may or may not include rounding.
- Q2: Flat dither** - In this case, we have $F(z) = 0$.
- Q3: Noise shaped** - $F(z)$ is some non-zero filter function and $DITHER = 0$.
- Q4: Dithered and shaped** - Both $DITHER$ and $F(z)$ are non-zero.

There are various reasons for choosing one or the other of these schemes, depending on the application. In any case, the least acceptable seems to be pure truncation. In Q1 and Q3, the amount of energy in the error is the same. In Q3, filtering is used to shape the spectrum of the quantization error, but it cannot reduce the absolute energy of the error. In Q2 and Q4, dither noise is added that raises the absolute noise floor by 3 dB. In Q2, the error will be white. These are not all of the possible arrangements, but they are the most commonly discussed ones [1-4].

The point of dithering is that the addition of some noise breaks the correlation between the quantization error and the signal. UPDF makes the first moment of the error uncorrelated, and TPDF makes the first and second moments uncorrelated. Obviously, other distribution functions could be used to decorrelate higher moments of the error, but these will not be discussed here.

AND WHAT SPECTRUM IS THAT ANYWAY?

The point of noise shaping, as mentioned above, is to move the noise energy into a portion of the audio spectrum where it is less audible. Following this argument, one would then conclude that the best spectrum for the noise is one that places the amplitude of the noise just below the threshold of hearing at each frequency. One need only look at the curves for the threshold of audibility and design a filter, $F(z)$, that produces exactly this curve. This is where the trouble begins.

IT ALL STARTED WITH FLETCHER-MUNSON

In 1932, a pair of researchers at Bell Laboratories decided to clear up what had up to then been somewhat of a mystery: how do we perceive the loudness of sinusoids at different frequencies and levels? For their experiment, they used vacuum-tube oscillators and amplifiers to produce tones that were presented through a single speaker [5]. Subjects were positioned directly in front of the speaker. They were then asked to match the loudness of tones at different frequencies. The result of these experiments was the ubiquitous set of curves shown in Figure 2. This is a copy of the figure in the original article.

One might ask whether the equipment available to them in 1932 was adequate to the task. They describe the harmonic distortion of the signal as follows:

"The receivers were of the electrodynamic type and were found to produce overtones of the order of 50 decibels below the fundamental. At the very high levels, distortion from the filters was greater than from the receivers, but in all cases the loudness level of any overtone was 20 decibels or more below that of the fundamental. Experience with complex tones has shown that under these conditions the contribution of the overtones to the total loudness is insignificant."

Although their conclusion is probably correct (that the overtones did not affect the final result), that level of distortion in a psychoacoustic experiment would not be considered acceptable today.

Now let us turn to the set of curves shown in Figure 3. These are from Stevens and Davis [6]. Note that they differ in many places from the Fletcher-Munson curves *by more than 10 dB*. These curves were taken using earphones only. Note that all the complicated behavior in the 5-to-10 kHz region disappears. We must conclude from this that the non-monotonic behavior above 5 kHz is caused by

the acoustical environment, and is not directly a product of the mechanism of loudness perception of the inner ear.

There have been numerous further studies using more and more refined equipment, but they still display large variations based on speaker position and acoustical environment. We must conclude that *there is no absolute threshold-of-hearing curve that is independent of the way the sound gets to the inner ear*. Even such simple operations as presenting the sound through two speakers placed to the left and right of center, as you might find in a living-room environment, is sufficient to change the curves by more than 6 dB.

Note also that the shape of the curves changes greatly from one loudness level to another. In the studio, we are required to choose one curve to design our noise-shaping filter with. This means we must decide not only what the acoustical environment (speaker/listener positioning, reverberation) of the ultimate presentation will be, but also on exactly what the listening level will be. Thus we conclude that *there is no absolute threshold-of-hearing curve that is independent of absolute level*. Spectrally shaped noise that is inaudible at one level will be audible if the level is increased, and the resulting perceived loudness of the noise will depend on the exact overall amplitude and on the spectral shape in relatively complicated ways.

Given all of this, how can we decide how to design our noise shaping filter? The short answer is, of course, that we can not make that decision, since we cannot control how the sound will ultimately be presented. That notwithstanding, CDs are being issued today with noise shaping, and most people agree that there is some advantage to using noise shaping in most cases.

Even though there is some uncertainty in the exact shape of the threshold-of-hearing curves, we can all agree that some general trends are observed, which are that the ear is the most sensitive in the range of frequencies from maybe 1 kHz to 5 kHz. It is increasingly less sensitive in the high frequency range and increasingly less sensitive in the low frequency range. Since there are relatively few critical bands in the high range, it makes sense to place most of the noise energy in the high frequencies. Any noise shaping filter that does this is defensible. Figure 4 shows an example of a range three different noise-shaping filters that are all perfectly acceptable for the reasons given above. The only way to decide among various filters that accomplish this is to listen to them and judge which one sounds better. In modern CD production, judgments like this are generally the responsibility of the mastering engineer.

Note that there is one other technical consideration in choosing a noise-shaping filter, and that is the possible loss of dynamic range from using higher-order

noise shaping. Figure 9 shows five examples of dithered quantization of increasing order. Note that the apparent amplitude of the higher-order shaping is considerably higher than the low-order (or flat) shaping. In general the absolute level of the quantization error is so low that this difference is not significant, but it does place a practical upper limit on the strength of the noise-shaping filter. Note also that some audio equipment reacts adversely to the presence of strong high-frequency energy. For the noise-shaping to be effective, the entire audio chain through the final speaker must be capable of accurately reproducing the quantization energy as well as the signal. Any nonlinearities in the system violate the assumptions involved in using noise-shaped dither in the first place.

WHY WOULD YOU NOT USE NOISE-SHAPING?

Given the "obvious" benefits of noise shaping in general, why would we ever *not* want to use it? Again, the answer to this depends heavily on what the ultimate use of the signal is. If we are just talking about recordings of music that will be played at home, in the car, or over the radio, then there is no reason not to use noise-shaping, and the aforementioned advantages of using noise-shaping. If, however, your sound is expected to be modified before use, then the situation is different. There are many thousands of CDs produced each year that contain sound effects or libraries of musical instrument tones for sampling synthesizers. For both sound effects and sampler libraries, it is expected that the sound will have vari-speed applied to it, *i.e.*, it will undergo a sampling-rate change to perform the sound at a different speed and pitch. If we change the playback speed of a sound with noise-shaped dither, *all* the spectral components of the signal are shifted, *including the dither noise*. Figures 10 and 11 show the spectrum of a signal with noise-shaped dither before and after the vari-speed operation has been done. Note that the noise energy has been shifted down right into what is generally agreed to be the most sensitive region of human hearing. Although the specific amount of vari-speed applied in this case was highly exaggerated for maximum impact, *any* amount of vari-speed violates the assumptions used to compute the noise shaping filter. Thus we must conclude that *sounds that are destined for sound-effects libraries and sampler libraries must have either no dither (Q1) or flat dither (Q2)*. They *must not* have any kind of noise-shaping applied.

WHY WOULD YOU NOT USE DITHER?

There are commercially-available quantization devices that do not inject any dither noise (choice Q3 above), but rely entirely on the truncation error to drive the noise shaper. The argument usually used to defend this choice is that it does not *add* anything to the signal, and it does not raise the noise floor. If the source

of the signal is a high-resolution converter of some kind, one can argue that the truncation error is relatively wide-band and is perfectly adequate to drive the noise-shaping process. There are several problems with this approach:

- (1) The truncation error of the signal is highly correlated with the signal, so it will be modulated by the signal. This produces a "pumping" noise floor that will be louder and quieter depending on the music. It is then impossible to claim that the quantization error is uniformly below the threshold of hearing since sometimes it will be much louder than at other times.
- (2) When the signal dies away, and even if it achieves zero values, the noise shaping filter will settle to some limit cycle and will continue to oscillate forever at some fixed frequency. Since all the limit cycle energy will be concentrated into a single harmonic series, it can be quite audible.

For these reasons, we conclude that *it is not reasonable to have a quantization system that does not add dither noise of some kind*. The problems of doing so far outweigh the supposed advantage of having the minimum error energy.

AND WHAT ABOUT DATA COMPRESSION?

There are some commercially-available products on the market today that use data compression to reduce the bandwidth of the audio signal. All the ones in current use for high-quality reproduction are some kind of sub-band encoding, where the signal is split into some number of frequency ranges and each range is coded by a certain number of bits. The differences among the algorithms have to do with how signal is broken into bands and how the number of bits are "spent" among the various bands. One might ask if there is an interaction between the method of quantization and the success of the data compression algorithm. A more practical way of phrasing the question is to ask if there are quantization methods that interact better with data compression.

Figures 12 through 15 show some of the results from a single widely-available commercial data reduction technique. Figure 12 shows how with TPDF the algorithm "reduces" away the high-frequency band. Figure 13 shows that high-order shaped noise can excite an instability in this algorithm causing it to alternately reproduce and eliminate the high-frequency band, producing extreme "pumping" of the high-frequency energy. Although not all data reduction algorithms show this pathological behavior, we must conclude that *one cannot blindly assume that quantization for linear PCM can be used for data-reduction with impunity*. One must always test to see if the results are going to be acceptable.

In general, it should be clear from this demonstration that *sound destined for a medium that uses data compression should be dithered differently from a sound that is destined for a Compact Disc*. This leads us to the conclusion that the mastering engineer should ideally produce several different versions of each album for each different release format. The difficulty of actually doing this combined with the cost of doing this probably means that it will not be done, and the results of Figure 13 will probably be commonplace.

Note that all data reduction algorithms have the possibility of accepting high-resolution material rather than sound that has already been quantized to 16 bits (although the physical systems that are available today to do the data reduction prior to mastering generally take only 16 bit signals as input). It would seem that there would be some advantage to delivering an unquantized master for release on compressed media and a quantized master for CD production. It would require, of course, that the encoding devices accept high-resolution data (20 or 24 bits). This is equivalent to saying that one will use the error inherent in the data compression and decompression process to accomplish what we do with dithering and noise shaping. This is not an unreasonable position to take, but its success depends strongly on the exact details of the data compression and decompression algorithm.

STARTING AND STOPPING

It has become common practice in CD mastering to require that the sound between tracks on an album be forced to all zeros. This means that any dithering on the sound must be turned off before the gap between tracks and turned back on afterwards. There are several ways to do this:

- (1) If UPDF is used (and no rounding is used in the quantization), then it will happen automatically, since the truncation will produce all zeros when the signal is zero.
- (2) With any other dithering algorithm, the dither must be turned off explicitly at some point.

The problem with (1) is that it is dependent on the signal. If you have, say, a long fade-out, then the dither will appear to come and go with the signal. This has the effect of re-introducing this modulation between the signal and the dither that we have been trying to get rid of.

The problems with (2) are several.

- (a) Since the quantization energy is constant, there is no way to fade it out gradually. It can only be on or off. This means that when you turn it off, there will be a "thump." There is no way to eliminate this.
- (b) With noise-shaping, the shaping filter always has some state and some limit cycle. Even if the signal is gone and the dither is set to zero, the filter itself will continue to "sing" forever. This means that if you decide to turn off the dither, you must clear the "state" of the noise-shaper as well.

Clearly, the most logical choice would be to use TPDF with noise shaping (Q4) and leave the dither and noise shaper running in the inter-track gap. This would require a change in the way CDs are mastered today, but would yield the smoothest transition from one cut to another.

AND WHAT ABOUT EDITING?

If noise shaping is used, then the ultimate quantization noise does have a non-impulsive autocorrelation that is due to the shaping filter. This means that if you overlap two or more signals that have already been quantized, it is possible to get interactions between the quantization errors of the different signals. For this reason, we recommend that all editing be finished before the quantization algorithm is applied. This requires that the quantization be the *last* step before the CD pressing master is made.

IN SUMMARY

We conclude with the following recommendations:

- Flat, triangular dither (Q2) should be used for sounds that will have any kind of vari-speed or sample-rate conversion applied, such as sound-effects libraries or sampler libraries.
- Noise-shaped dithering (Q4) should be used for 16-bit linear PCM release formats. The exact spectral shape of the filter can vary widely and still produce acceptable results. The mastering engineer should be provided with ways to change or customize the spectral shape depending on the application.
- Signals destined for media that incorporate data reduction should be left at high resolution and completely unquantized before compression. Data reduction algorithms should make use of high-resolution input data and be designed so that the resulting error performs the same function as dithering and noise shaping (*i.e.*, reduces the audibility of the processing error).
- Equipment manufacturers should provide the users with a wide variety of dithering and quantization options so that the result can be matched to the application in a flexible manner.

We hope that these observations will serve as a basis for future discussion of the role of quantization in the music mastering process.

IT FIGURES:




Figure 5: Original piano music. This example is about 22 seconds long.




Figure 6: Original piano with fade-out of 3 dB/second at full resolution (no truncation). Note that the tones in the last half are not visible at this magnification.



Figure 7: Piano with truncation. Note that the last few notes are well below the least significant bit.




Figure 8: Piano with triangular dither of amplitude twice the LSB. Note that only the peaks of the first few notes are visible in the waveform. The last few notes of the piano are clearly audible although they are well below the level of the LSB.




Figure 9: Four copies of the example with (1) flat triangular dither, (2) 2nd-order dither, (3) high-order dither weighted towards low-frequencies, (4) high-order dither with middle weighting, and (5) high-order dither with high-frequency weighting. The energy of the dither in all 5 cases is the same. The increase in amplitude is due to the noise shaping.




Figure 10: Spectrum of a signal with noise-shaped dither. Note that the partials of the signal are clearly visible in the low frequency range (the frequency scale is linear, 0 to 24000 Hz).




Figure 11: Spectrum of a signal with noise-shaped dither that has had varispeed (sample-rate conversion) applied. Note that the high-frequency energy of the spectral shaping is now shifted down to a region of the spectrum that is highly audible (the scale is linear, 0 to 24000 Hz).




Figure 12: Spectrum of a signal with flat, triangular dither that has been processed by a widely-available commercial data reduction technique. Note that the technique has "reduced" away all the information, dither or otherwise, in the region above about 15 kHz.




Figure 13: Spectrum of a signal with noise-shaped dither that has been processed by a widely-available commercial data reduction technique. The great amplitude of the dither has excited an instability in the data reduction algorithm. The amplitude of the entire high-frequency band is oscillating between zero and full amplitude. The time scale represented is about 1/2 second.




Figure 14: Spectrum of the previous signal during the portion when high frequencies are reproduced. Note that the general spectral shape is preserved, but the algorithm has "spent" a great deal of its bandwidth allocating reproducing the exact dither curve.




Figure 15: Spectrum of the above signal during the portion when high frequencies are not reproduced. Again, the material above 15 kHz is missing (this spectrum looks different because it involves fewer points in time than the previous spectrum).

REFERENCES:

- [1] M.A. Gerzon and P.G. Craven,, "Optimal Noise Shaping and Dither of Digital Signals," presented at the 87th Convention of the Audio Engineering Society, J. Audio Eng. Soc. (Abstracts), vol. 37, p. 1072, Dec. 1989, preprint 2822
- [2] R.A. Wannamaker, "Psycho-Acoustically Optimal Noise Shaping," presented at the 89th Convention of the Audio Engineering Society, J. Audio Eng. Soc. (Abstracts), vol. 38, p. 871, Nov. 1990, preprint 2965
- [3] S.P. Lipshitz, J. Vanderkooy, and R.A. Wannamaker, "Minimally Audible Noise Shaping," J. Audio Engineering Society, Vol. 39, No. 11, pp. 836-852, Nov. 1991
- [4] M. Akune, R.M. Heddle, K. Akagiri, "Super Bit Mapping: Psychoacoustically Optimized Digital Recording," presented at the 93rd Convention of the Audio Engineering Society, San Francisco, Oct. 1992, preprint 3371
- [5] H. Fletcher and W.A. Munson, "Loudness, Its Definition, Measurement and Calculation," J. Acoustical Society of America, Vol. 5, pp82-108, October, 1933.
- [6] S.S. Stevens and H. Davis, "Hearing - Its Psychology and Psysiology," New York: Wiley, 1938, p. 124

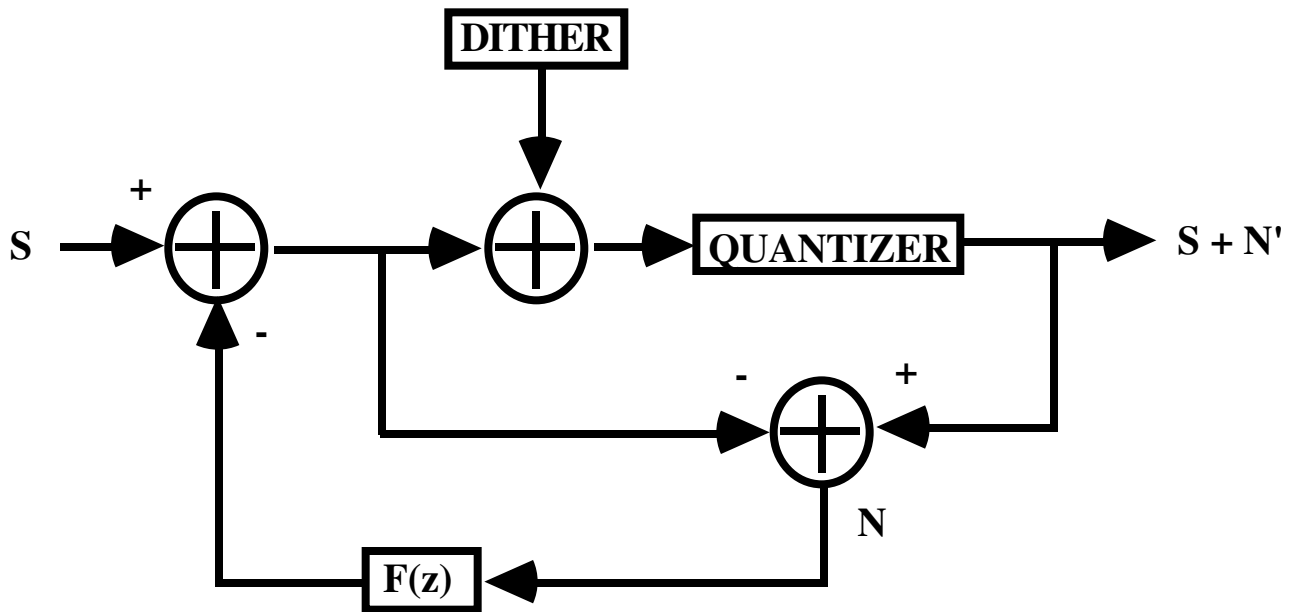


Figure 1: General block diagram of noise-shaped dithering. For unshaped dither, $F(z)$ would be zero. For undithered noise shaping, the dither would be zero.




Figure 2: The equal-loudness curves reproduced directly from the 1933 article of Fletcher and Munson [5]. The signal used for these measurements were presented from a single loudspeaker placed directly in front of the subject.

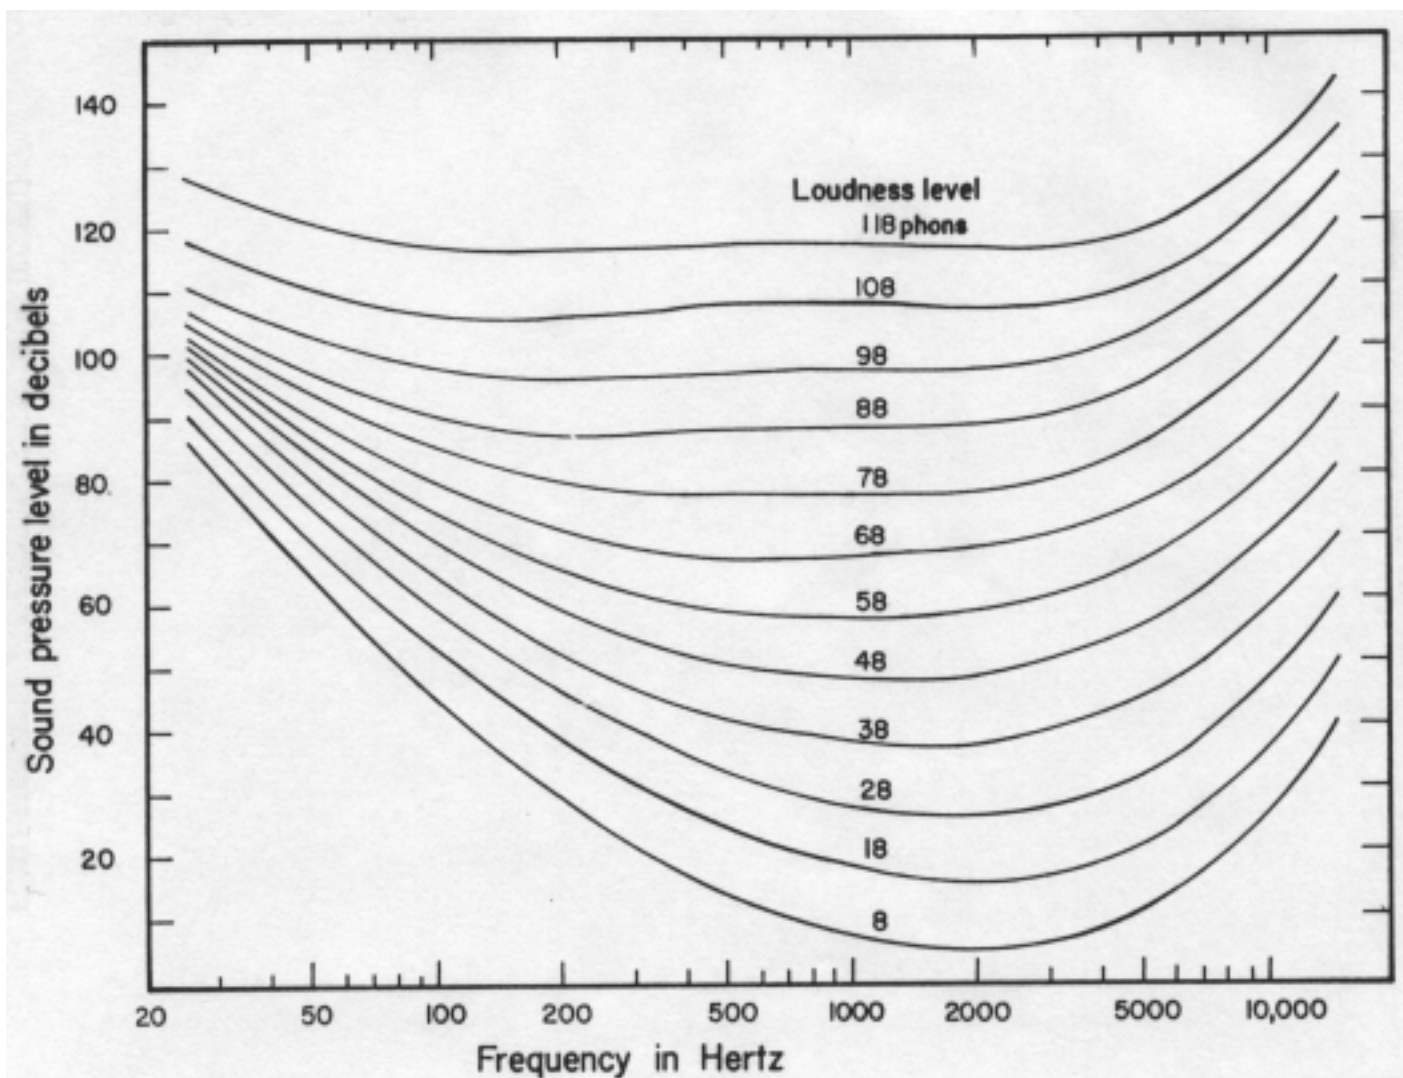


Figure 3: Another set of equal-loudness curves from Stevens and Davis [6]. These were derived using earphones.




Figure 4: A set of three plausible curves for the frequency response of a noise-shaping filter. There is no absolute way of deciding which one of these curves is "best." The decision must be a subjective one.

