Trevor Owen de Clercq
October 5th, 1998

Chapter 3 Abstract

In this chapter, John Watkinson finally starts off into the basic mechanics of audio digitization and vice versa. He refers to this process as "conversion" (the title of his chapter), by which he means the conversion of analog signals into the digital domain and digital code into the analog domain. Some of the concepts he describes, such as aliasing and the Nyquist limit, I've already discussed in other reports; the topic of dither, however, is basically new to me and thus the one on which I focused this week.

To truly understand the nature of dither, I also referenced two other texts, one being the excellent introduction to digital audio by Ken Pohlmann entitled Principles of Digital Audio and the other being Erik Sohetina's Complete Guide to Digital Audio Tape Recorders. These two texts seemed to give fair laymen's explanations of the dithering concept. In essence, dither is the addition of noise to an analog signal to prepare it for quantizing. At first, it may seem counterintuitive to *add* noise to a signal since the goal of audio is to eliminate as much noise as possible, but as will be seen, this extra noise, when combined with noise shaping, actually increases the SNR (signal-to-noise ratio) of the digital system. The purpose of dither is to deal with soundwaves whose amplitudes fall within one quantization level. In a sense, therefore, dither is kind of related to anti-aliasing; anti-aliasing, of course, deals with soundwaves whose frequencies fall within one sampling rate. Thus, the analogy can be made: dither is to amplitude as anti-aliasing is to frequency. Unfortunately, unlike aliasing, soundwaves with minutely varying amplitudes cannot be filtered out since they lie within the audible sonic spectrum.

The problems incurred by soundwaves within one quanztization level are particularly annoying to the human ear. Since these soundwaves fall with one bit, it will appear to the quantizer as if they were DC (direct current). Apparently, this direct current representation in digital form becomes very noticeable as distortion (not noise) once the data is transferred back to the analog domain. To combat this distortion, dithering is added. Noise, as it would seem to be implied, is a lesser evil than digital distortion.

The mechanics of dither are simple. Dither is simply a sort of code for the signal it is representing. By adding random noise to the signal, one can detect the true amplitude of the signal by plotting the oscillations of the dithered (noisy) signal between the two quantization levels. A signal that is closer to one level will statistically reach that level more than a signal that is closer to the other level. For example: say two bits represent the voltages 1 and 2 and an undithered signal is wavering around 1.75 volts. By adding random voltages between 0 and 1 to the signal, the dithered signal will more often peak above 2 volts than fall below 1 volt. Similarly, a dithered 1.95 signal will peak even more often above 2 volts than the 1.75 volt dithered signal. It all comes down to statistics and probabilities. By averaging muliple generations of the dithereing process, it is possible to come extremely close to the original waveform that was unrepresentable prior to the dithering process. It is my understanding that this averaging process is what is referred to as "noise shaping".