Dithering

Introduction

On Internet I have found a very nice document and small program to enhance recordings after editing of recorded .WAV-files. Stuart Binder of Solution Technologies is the author of it. This is of great importance for people who edit their own recordings, so that I asked Stuart to permit for publication on this website.
Mind however that good editing programs as Adobe Edition include dithering facilities! In case you work with such a program, Stuarts lecture is instructive at least.

Stuart:

The utility: Dither is for dithering 24 bit stereo or mono WAV files to 16 bits for CD mastering. The program will dither from 24 to 16 bits, or from 16 bits to 8 bits.

What is dithering ?

The best way to explain dithering is to first examine what happens when bit depth is reduced without dithering. The alternative to dithering is to simply truncate (or round off) the least significant bits. Since these bits control the quietest parts of the signal, they can be removed without drastically damage-ing the recorded material. Most software that changes bit depth works this way. The problem, however, is that truncation does subtly damage the sound – the difference between a truncated recording and a dithered recording is obvious if you know what to listen for (more on this later).

Dithering allows the least significant bits of recorded material to be removed, while retaining some of the information encoded in those bits. It is the best way to reduce a 24 bit master recording to 16 bits for CD format. Dithering is accomplished (believe it or not) by adding a small amount of noise to the material before removing the least significant bitls.

The best explanation I have heard as to why this works involves a visual analogy. Hold your fingers in front of a page of print, and try to read the print. It will be difficult (or impossible), since your fingers hide some of the words. The page is being quantized by your fingers in much the same way that an analog audio signal is quantized by an A-D converter – and the same way that a 24 bit signal is further quantized by reduction to 16 bits. You can still see the print, but the resolution is not good enough to read the print. This is truncation (part of the data on page is truncated by your fingers). Now move your hand quickly from side to side, and you'll find that you can read the print without any trouble. Your fingers are still between you and the page, and they are no smaller. Yet it is now possible to read the print. This is because some noise (the motion of your fingers) has been introduced into the process of quantization. The threshold between what is visible and what gets cut off is no longer stationary, so that while none of the print is visible at all times, all of the print is visible some of the time.

Audio dithering makes use of the same principle. The threshold is kept in motion between the part of the signal that is retained and the part that is discarded, so that some of the discarded material is still audible in the final product.

Hear for Yourself
Dithering to 8 bits can be very instructive, as it is like listening to a 16 bit (CD quality) product under a 256x "microscope." If you have any doubts at all about the effectiveness of dithering, the following experiment will prove very enlightening. The difference is clearly audible – even on cheap multimedia speakers!

Create a 20 second 16 bit WAV file with a 500Hz sine wave. Use your WAV editor to fade the file out over its entire length – so that it starts at 0dB and ends at -infinity. Save the file, and then truncate it to 8 bits. When you listen to the result, you'll know why nobody works in 8 bits! The 500Hz sine wave is still there, but with only 8 bits to represent it, the "stair stepping" is clearly audible – you can hear artifacts – overtones that shouldn't be there (this is, after all, a sine wave). As the signal fades, it deteriorates even more – it sounds awful.

Now take the same 16 bit source file, and dither it (use the default – TPDF Dither with Noise Shaping enabled). When you listen to the result, you may well be amazed. The dithering process provides a soft, low-level noise floor to the signal, and the sine wave is clearly audible above this noise floor. It sounds like a real sine wave too – all the way to the end of the test. The default options should work well for most source material, but try experi-menting with the other dithering options to see what you like best.

How to use the program: Dither?

Dither may be launched by double-clicking its icon, or by dropping a file (or files) to be dithered on its

icon. Dither's main window is a dialog for collecting information about the files to process and the processing methods.

Files to Process

The Select button (to the right of the "Files to Process" area) allows selection of a file or files for processing. Any files selected in this way will replace the previous list in the "Files to Process" area.

Processing Method

Truncate – does just what it says – it discards the 8 least significant bits.
Round – gives the LSB a vote (well, half a vote, anyway) in the final outcome. The results of rounding are not generally much better than truncating.
White Dither – This option is where things really start to get interesting. White Dither uses random (white) noise to dither the material. It sounds much better than truncating or rounding, but you may still hear some artifacts above the noise floor.
TPDF Dither – This option also uses white noise, but the amplitude of the noise is weighted toward the center. This type of noise usually produces better results than white noise. TPDF stands for Triangular Probability Density Function.
Noise Shaping – This option is available for use with White Dither and TPDF Dither. It will have no effect when truncating or rounding. The Noise Shaping option pushes as much of the noise as possible into the higher frequencies, where it is less audible. The result is a quieter noise floor.

Dithering Notes

Dithering should be the absolute last thing you do to an audio file. Any processing done after dithering (including something as simple as changing the levels) will destroy the effect of the dithering. Do all of your processing (levels, eq, fades, etc.) in 24 bits. Then dither the file to 16 bits as the final step before burning a CD.

In this first release (ver. 1.0ß) the dithered file will be placed in the same directory as the original, with "_d" appended to the filename.

Please let me know what you think of Dither.
Email your comments to sbinder@stdesign.com.

Thanks!

Stuart Binder
Solution Technologies