A Cost-Effective Sound Expander

Lloyd Watts

March 15, 1984

Bachelor's Thesis
Engineering Physics
Queen's University
Kingston, Ontario
Canada

Abstract

We present a digital signal processing method and implementation of an audio time-stretching algorithm. The intention is to slow down speech and/or music by a factor of two without changing the pitch of the signal. The approach is to use a half-speed tape recorder to play back audio at 1/2 the original speed, and digitize the slowed-down audio with an analog-to-digital converter (A/D), apply a pitch-doubling algorithm to restore the original signal pitches, and then convert the pitch-corrected signal back to analog audio using a digital-to-analog converter (D/A). The pitch-doubling algorithm is implemented as a circular memory buffer, where the write head writes into the buffer at the sampling rate, and the read head reads from the buffer at double the sampling rate. This approach effectively doubles the sampling rate of the output signal, thus doubling the pitch frequencies of the slowed-down audio signal, restoring them to their original pitches. The algorithm can result in waveform discontinuities when the read head overtakes the write head, and thus the time-dilated signal can have some noisy artifacts. The system was built using a half-speed reel-to-reel tape recorder and an Apple II computer with a 1 Megahertz 6502 microprocessor with a custom A/D and D/A card built by the author. The pitch-doubling algorithm was implemented in assembly language in 10 lines of code.