

Class Overview

Introduction

There are typically four steps in producing a CD or movie soundtrack, as shown in Figure . In *tracking* sounds are recorded or synthesized and arranged in tracks. The tracks are then processed in the *mixing* stage to form a stereo or multichannel mix. The idea is to arrange the sounds spatially and spectrally, to manipulate their character for artistic purposes, and also to fix problems in the tracks. In *mastering*, subtle adjustments and fixes are made to the mix, and often its dynamic range is limited in preparation for *encoding* and printing on the target medium.

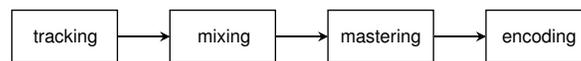


Figure 1: *Audio Production Process*

This class is about audio effects processing, and, in particular, how to build digital versions of the mainline audio effects used in the mixing and mastering stages of audio production. There are four categories of processing commonly employed in mixing and mastering:

- Dynamic Range Control,
- Reverberation,
- Equalization and
- Spatialization.

Other specialized processing, such as delay effects (including echo, chorus, flanging and phasing), distortion, pitch and time stretching and noise removal, is also used. In this class we will explore dynamic range control, reverberation, equalization and distortion, and in homework and laboratory exercises you will build examples of each. We will also touch on some specialized processors; you may wish to choose one to study as your project.

Dynamic Range Control

Dynamic range control refers to the manipulation of the level of a signal. This may be desirable, for instance, in the case of a singer who moves closer to and further from the microphone while singing, causing the level to unintentionally rise and fall. It is commonly used to make current pop CDs quite loud, despite their 16-bit integer samples.

As illustrated in Figure , dynamic range control may be accomplished using a feed forward architecture in which the signal level is detected and used to compute a gain, which is applied to the input. The gain is computed based on a desired output level as a function of input level.

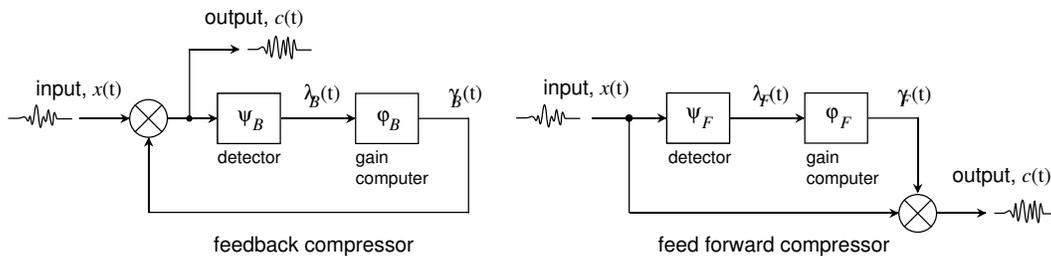


Figure 2: *Feed Forward and Feedback Compressor Architectures.*

Different applications are developed by choice of gain computer and detector. A gain computer which reduces the dynamic range of louder signals results in a *compressor*, and may be used to make a bass or drum track more even. By designing the gain computer to impose a predefined maximum output level, the input is *limited*. A *noise gate*, which eliminates any low-level background noise appearing between notes in the input track, can be implemented by using a gain computer which takes the signal level to zero when the input is small.

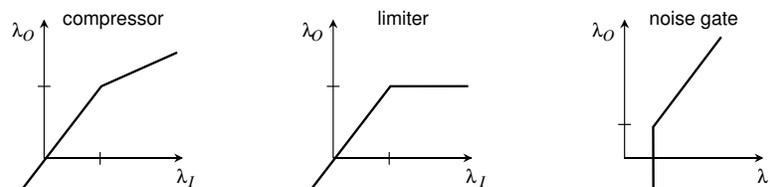


Figure 3: *Compressor, Limiter and Noise Gate Input Level-Output Level Relationships.*

If the detector determines level by examining the input signal over a long time period, the estimated level will change only slowly, and the output will be “transparent”—that is, it will sound much like the input. If the detector only considers a short period of time in evaluating signal level, it can change the envelope of individual notes and the

character of the track. Finally, the detector (or the entire architecture, for that matter) can be applied to selected frequency bands, for instance, to reduce loud sibilence.

The dynamic range control unit begins by exploring the notion of signal level and introducing common processing architectures. Signal level detection and gain computation are covered, followed by an overview of applications such as limiting, gating and de-essing, and a presentation of improvements including efficient FIR detection and aliasing elimination. The problem sets will focus on details of detection and gain computation, with the laboratory exercise being to modify a simple compressor.

Reverberation and Room Acoustics

The acoustics of a space can significantly contribute to the feel of a piece: imagine the “sound” of a jazz hall, a dungeon, outdoors. For this reason it is desired to be able to artificially add environmental cues or *reverberation* to tracks or a mix.

In addition, the tracks of a mix are not often recorded under the same acoustic conditions; consider vocal booths and drum rooms. As a result, artificial reverberation is commonly added to a mix so as to make different tracks feel as if they belong together.

There are two approaches to implementing artificial reverberation in common use today. In one approach, delayed copies of the signal are filtered, mixed and fed back to delay line inputs. The idea is that the process in some ways mimics what happens in actual acoustic spaces, with reflecting surfaces and objects filtering impinging wavefronts and redirecting them to other reflecting surfaces and objects. Loosely speaking, such artificial reverberators are called *feedback delay networks*.

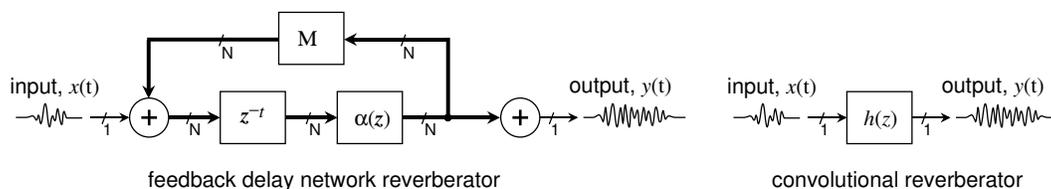


Figure 4: *Feedback Delay Network and Convolutional Reverberators.*

In the other approach, it is assumed that the acoustic space is approximately linear and time invariant, and can be characterized by its impulse response. Artificial reverberation may then be applied to an input signal simply by convolving it with the desired impulse response. In this approach, impulse responses can be measured and manipulated or synthesized to achieve a desired artistic effect. As an example, Figure shows the time evolution of the a balloon pop recorded in Memorial Church.

In the reverberation unit, the focus is on convolutional reverberation. The unit starts with room impulse response measurement, presenting Golay code and allpass chirp techniques. Reverberation acoustics is then studied, including the image method and Sabine

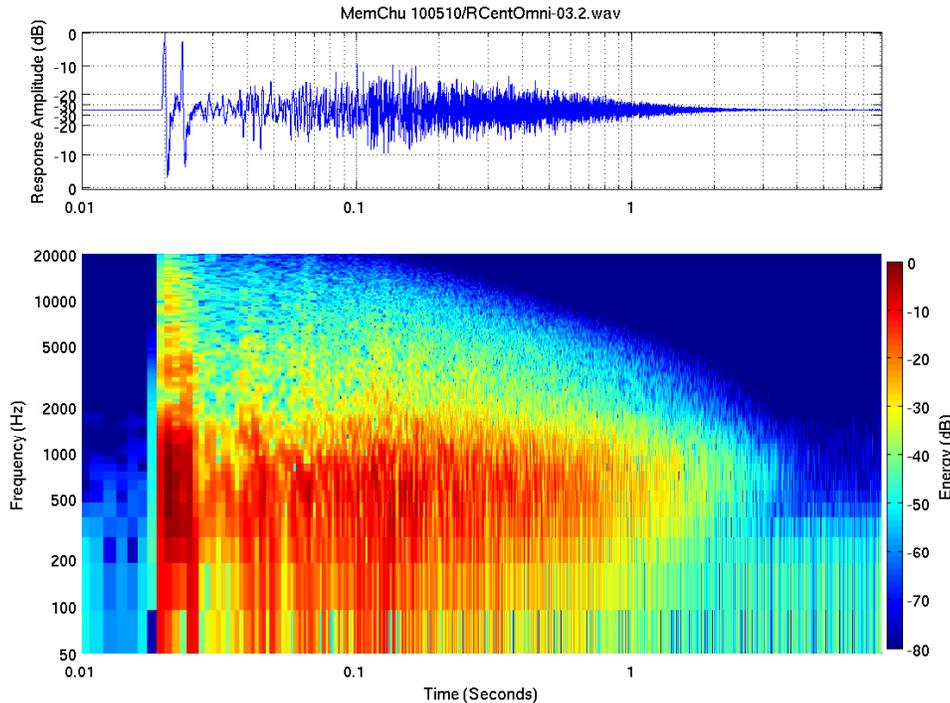


Figure 5: *Memorial Church Balloon Pop Response and STFT.*

theory. Impulse response analysis and synthesis techniques are covered, and reverberation psychoacoustics is discussed. Problem sets center around reverberation impulse response analysis and synthesis, and the laboratory exercise will be to measure, analyze and manipulate the impulse response of an acoustic space. A convolutional reverberator will be provided to listen to manipulated and synthesized impulse responses. A feedback delay network reverberator will also be provided, and you will learn how to modify it so as to match to sound of a given space or achieve a particular psychoacoustic effect.

Delay Effects and Distortion

There are a number of situations in which a track is purposefully distorted so as to lend a certain character to its sound. It's not unusual, for instance, to send a track in digital form out to an analog tape deck and back to give it a bit of "warmth." In a much less subtle example, guitar amps are often driven to saturation with pleasing results, and guitars tracks are many times recorded clean and distorted appropriately in mixing.

A common architecture for a distortion processor is the cascade of filtering and non-linear elements, as shown in Figure . The nonlinearity has the effect of increasing the bandwidth of its input, and to accommodate the wideband result, it is processed at a high sampling rate. In this portion of the class you will learn how nonlinearities alter the

spectral content of a signal, and about antialiasing filter design.

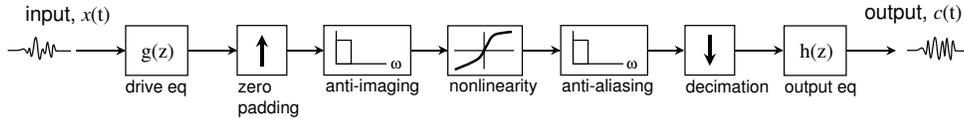


Figure 6: *Distortion Processor.*

Delay effects, including echo, chorus, flanging and phasing, are widely used on keyboards and guitar. They more or less add a series of time-delayed copies of the track to form their output, as shown in Figure 7. If the time delay is sufficiently large, the signal copies will be heard as distinct echoes or a chorus-like sound. When the delay between successive echoes is small enough that it is comparable to the period of a signal in the audio band, the process will be perceived spectrally. In this case a changing time delay results in a changing equalization, as seen in the example in Figure 7. In this portion of the class you will learn how to design fractional and time-varying delays and allpass filters, and will be given a flanger/phasor to modify.

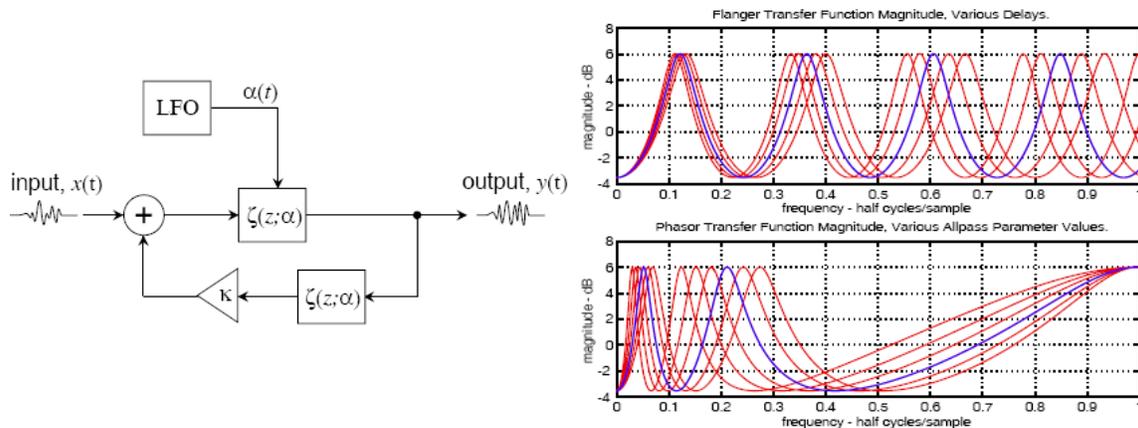


Figure 7: Flanger/Phasor Architecture, Example Transfer Functions.

Equalization and Filtering

Equalization and *filtering* are each the manipulation of the frequency content of a track, with filtering commonly referring to the removal of a band of frequencies and equalization often referring to a more subtle adjustment as a function of frequency.

In many cases equalization and filtering can be used to fix problems with a track. Tracks will sometimes have unwanted frequency components—say, a 60 Hz hum or a rumble from road noise—and filters may be used to remove them. There are other cases

where certain frequency components need to be enhanced; for example, a singer's lisp can be corrected somewhat by enhancing high frequencies.

One of the primary uses for equalization on tracks is to help separate different mix elements by having them occupy somewhat different frequency bands.

Equalization may also be used as an effect. Filtering the waveform to a band between 200 Hz and 3200 Hz (in combination with some other processing) makes the track sound as if it is being played through a telephone. In an architecture similar to that of the feed forward compressor above, the signal level may be detected and used to control a filter. In this way, for instance, the onset of a note can be made to have a very different timbre than its decay.

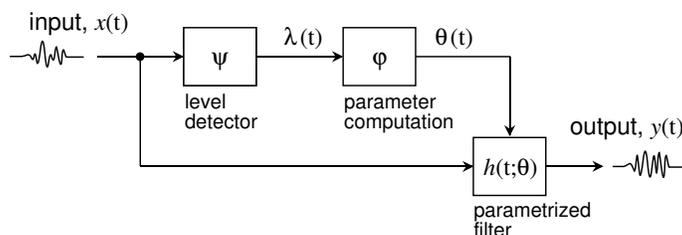


Figure 8: *Envelope Following Filter.*

It turns out that across a given genre, there is surprisingly little variation in overall spectral content. As a result, the mastering engineer will often (and probably subconsciously) use a program equalization which brings the power spectrum of the song in line with the standard.

Audio engineers have a preference for analog equalizers, for their transfer function characteristics, their controls and their signal path. The approach to equalization we present here is based on modeling analog equalizers to fix the transfer function and controls, and uses numerically robust filter structures so as to maintain signal integrity.

In the equalization and filtering unit, we first concentrate on peaking and shelving filters which are widely used in mixing and mastering, and which provide excellent building blocks for forming more complicated equalizers. Equalization psychoacoustics is then discussed, including presentation of the Bark and ERB frequency scales, and critical band smoothing. Linear phase and minimum phase filtering are covered, as are techniques for IIR filter design. The problem sets will cover IIR filter design techniques, and the laboratory exercise will include implementing a parametric section and maybe an envelope-following filter.

Course Topics

1. Dynamic Range Control (Compression)
 - Processing architectures
 - Detection and gain computation; analog detectors
 - Applications, architectures and improvements
2. Equalization
 - z-Plane, s-Plane and Fourier, Laplace relationships; bilinear transform
 - Canonical filters: parametric sections, shelf filters, cut filters
 - Time-varying and envelope filters
3. Distortion
 - Sampling rate conversion and antialiasing filter design
 - Distortion processing
4. Filter Design
 - Filter phase, linear and minimum phase conversion
 - Critical-band smoothing, Bark and ERB frequency scales
 - Frequency warping: warped FIR and IIR implementations
 - IIR filter design and Prony's method
5. Impulse Response Measurement
 - Golay codes, all-pass chirps, swept sinusoids
 - Weakly non-linear system measurement
6. Room acoustics
 - Room Impulse Response (RIR) analysis
 - Image method; Sabine theory
7. Artificial Reverberation
 - RIR synthesis
 - Low-latency convolution
 - Feedback Delay Network (FDN) reverberation
8. Modal Processing
 - Modal Reverberation
 - Modal effects (modulation, distortion, etc)