Subtractive Synthesis and Filtering Summary

Resources

Charles Dodge - Computer Music - Chapter 6 & 7
Miller Puckette - TATOEM - Chapter 8
Gordon Reid - Synth Secrets - Part 1 to 30
http://www.soundonsound.com/sos/allsynthsecrets.htm

Subtractive Synthesis

Subtractive Synthesis takes a rich/dense harmonic source and sends it into a group of one or more filters to shape the harmonics.

Typically used as an intuitive physical model where an instrument is broken into an excitation function and a resonator.

instrument	excitation function	resonator
trumpet	mouthpiece/lips	tubing & bell
voice	vocal cords	vocal tract
drum	stick/hand	drum head & body
guitar	string/pick	body
guitar	pick	string & body

One can often divide the resonator into a pitched resonator and a static resonator to model those parts of the instrument in which pitch changes.

In subtractive synthesis, pitch can reside in either the excitation function or the resonator or both.

In the case where pitch resides in the excitation function, it is typical that the resonator does not strongly shape the harmonics of this sound. That is, the filters which comprise the resonator will have a lower Q. This is because a higher Q filter will only let through harmonics near the filter frequency. An exception to this is when the resonator and excitation function both carry pitch information, so that the oscillator and filter are in a fixed frequency relation.

When pitch resides in the resonator, it is typical to use an excitation function which isn't pitched, that contains noise or an impulse. This is typical for modeling drums, percussive and plucked instruments. In some cases, real impulses are used (drum pads).

Filters

The basic types of filters are lowpass, bandpass, highpass and bandreject.

All of these filters can be derived from single or double cascaded allpass filters.

Highpass and lowpass filters have cutoff frequency, and slope.

A first order lowpass/highpass will have a slope of 6 dB/octave.

Bandpass/bandreject filters have two cutoff frequencies, slope and bandwidth (distance between cutoff frequencies); but are typically specified with center frequencies and quality or Q

Q = frequency/bandwidth so that a narrower bandwidth is a higher Q.

A second order bandpass/bandreject will have a slope of 3 dB/octave.

Higher slopes are acquired by cascading multiple filters.

A more detailed frequency response is acquired by placing many filters in parallel. (You could achieve the same frequency response by placing the same filters in series, but you would lose gain, and have a larger phase shift).

In some filters a higher Q is generated through feedback. In this case, the higher Q will cause a higher phase shift, and will break into oscillation at some points.

A typical and famous cascaded filter design is the Moog lowpass filter. This is comprised of four lowpass filters in series with a feedback section to increase the Q, this filter structure can move from lowpass to bandpass to an oscillator by simply increasing the Q control.

The state variable filter is a similar design, but with only two stages. In this case however, it is two integrators in a feedback arrangement which can give highpass, lowpass and bandpass simultaneously (with adjustable Q). Sometimes called a universal filter. When this filter oscillates, it becomes a multiphase oscillator.

Other filter techniques - IIR, FIR, biquad

More detailed filter responses can also be created with the FIR filter and the IIR filter.

Both of these filters are comprised of a series of multiple one sample delays, with the outputs of each delay multiplied by gain values (coefficients) and the results added together.

FIR

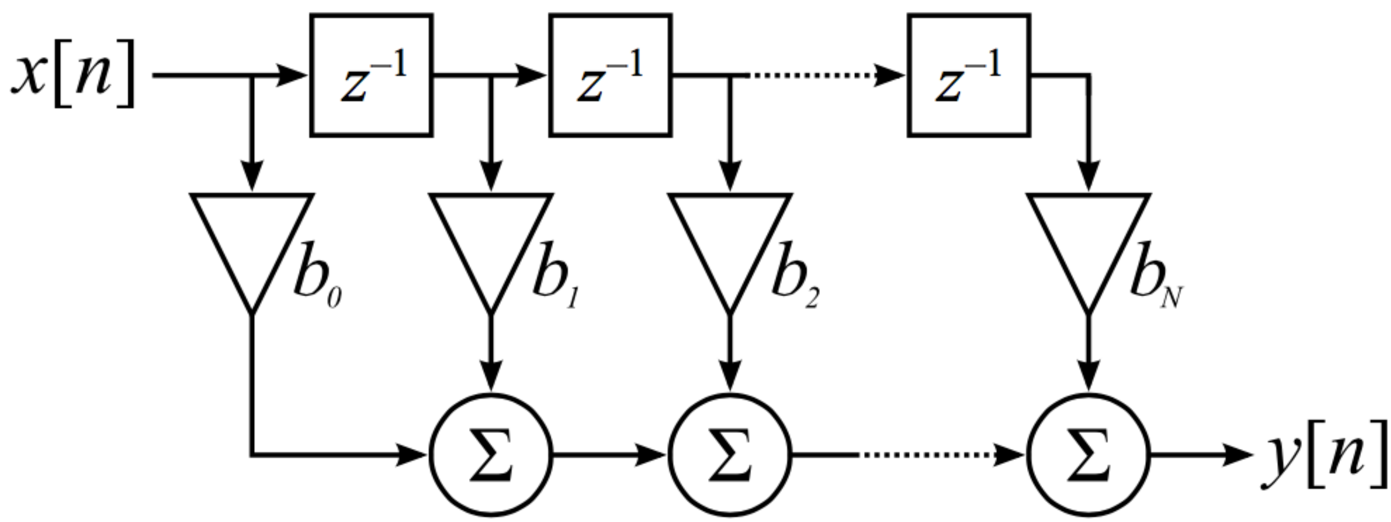

The FIR filter (meaning finite impulse response) has no feedback, and is also known as the convolution filter. For those of you with music tech experience, this is similar to a multi-tap delay.

The set of coefficients in an FIR filter is known as the impulse response. If one sends an impulse into an FIR filter, one gets the impulse response out.

The spectra of the impulse response is the same as the frequency response of the FIR filter using this impulse response. This gives you very easy control of phase and frequency response in an FIR filter.

FIR filters can get very inefficient. For a 20000 coefficient filter, one would need 20000 multiplies and adds for each sample. This would be 882,000,000 instructions per second - if stereo, 1700 MIPS.

For this reason, a spectral technique is often used to speed up FIR filtering.

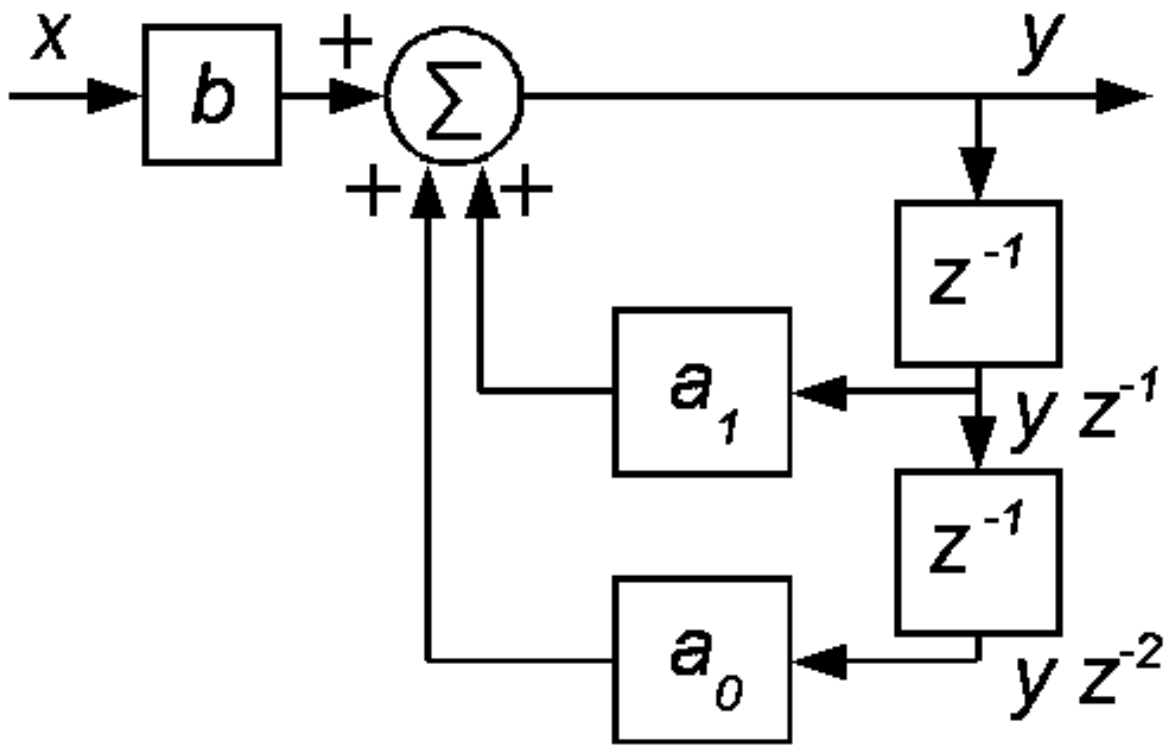

IIR

The IIR filter is similar in structure, but uses feedback. Because of this, each coefficient value applies to multiple delays (multiple times through the filter). Calculation of coefficients for an IIR filter is usually done with a filter design program which recursively fits the frequency and phase response required.

One technique to calculate IIR filters is LPC analysis.

IIR filters, because they use feedback, typically need less coefficients to achieve the same filter detail as an FIR filter.

Both filters require much computation to change the frequency response, so are typically not easy to change in real time.

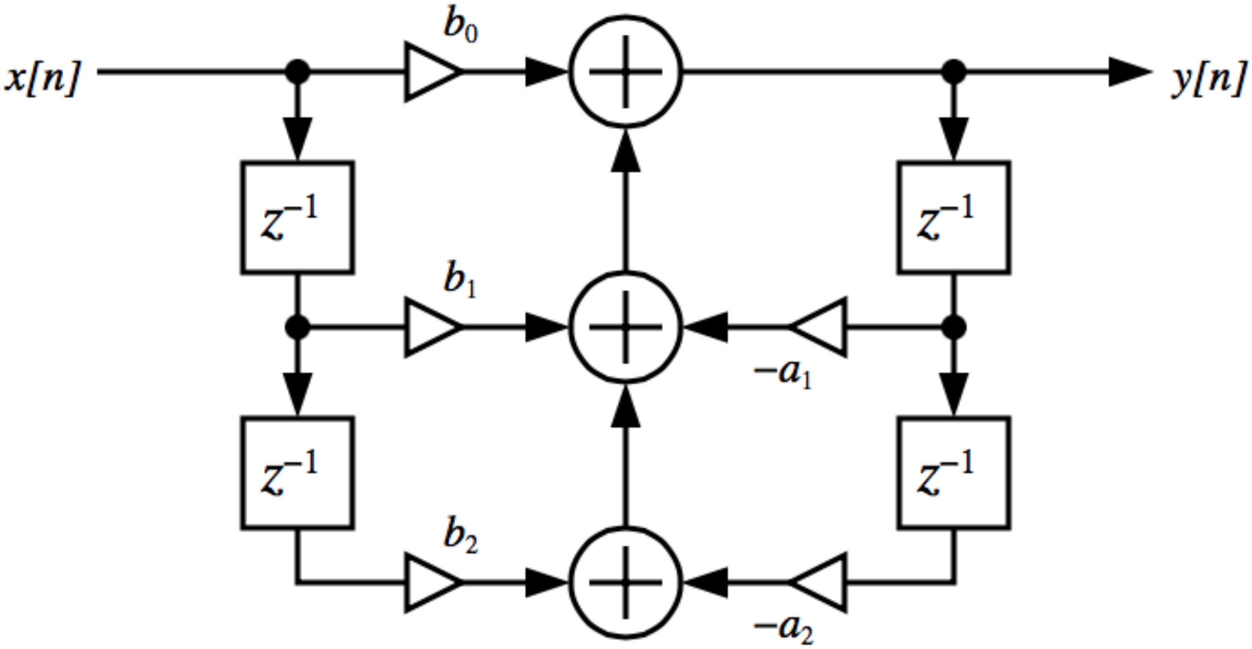

biquad

The biquad filter consists of a feedforward and a feedback stage connected together.

It uses 5 or 6 coefficients, and is a simple to use building block which can create all of the 4 basic filters (LP, HP, BP, BR) and more (shelf, peaking, butterworth, etc).

There are several commonly known techniques for generating these coefficients, check google for the RBJ cookbook, or musicdsp.org for the Zolzer biquad filter.