


"Types of Synthesis"

There are several types of sound synthesis used for generating instrument sounds in electronic music. They use methods ranging from combining basic waveforms to complex algorithms that mathematically reconstruct a real instrument's physical attributes. The following descriptions are simple introductions to the different types of synthesis, meant to help in understanding how the common types of sound synthesis work. The following synthesis methods are not exclusive. Many hardware and software synthesizers use a combination of these techniques to produce new, more realistic and/or interesting sounds.

Additive Synthesis

Additive Synthesis constructs sounds from the ground up by building a complex waveform from its basic elements. It works by mixing (summing) one or more simple waveforms, such as sine waves, together to create a more complex waveform. Below an illustration that shows how two sine waves can be summed to form a more complex waveform.



An example of additive synthesis.

This method of synthesis is theoretically capable of reproducing any sound. This is true because every sound can be broken down into a collection of sine waves. Unfortunately, natural sounds are extremely complex and require a great amount of processing to be recreated accurately.

Subtractive Synthesis

Subtractive Synthesis takes a reverse approach to generating sound. It starts out with a waveform rich with harmonics (such as a saw or square wave) and filters it to produce the desired output. This method can be used to effectively recreate natural instrument sounds as well as textured surreal sounds.

Granular Synthesis


Granular Synthesis uses a sequence of short sound grains to form a final longer output sound. Sound grains are relatively short segments of audio (usually from 10 to 100ms) that consist of frequency information in the form of a group of samples. The sound grain can be simple sine waves or complex sampled sound textures. The sound grains are played in a variable sequence with different variations on the amplitude, individual sound grain durations and total duration of the sequence of sound grains to form a final output sound/instrument. These variations define the output sound's envelope, pitch and length, while the contents of the sound grains define the output sound's "texture".

Forms of granular synthesis are commonly used to independently change the pitch/frequency or duration attributes of digital audio without affecting the other.

Amplitude Modulation (AM) Synthesis


Amplitude Modulation (AM) Synthesis is performed by combining two signals together. The source audio signal, the carrier, is multiplied by a unipolar modulation signal. A unipolar signal is simply a signal that contains only positive values (usually between 0 and 1). This process is typically used to alter the carrier signal in one of two ways.

The modulation signal can be used as an envelope which is applied to the carrier signal to determine the audio signal's amplitude over time.



AM synthesis used to apply an amplitude envelope.

The modulation signal can also be used to quickly cycle the carrier signal's amplitude to form two additional frequencies known as sidebands, forming harmonic or non-harmonic sounds.



AM synthesis used to add sidebands.

Ring Modulation (RM) Synthesis

Ring Modulation (RM) Synthesis is almost identical to AM synthesis (described above) with the exception that it uses a bipolar modulation signal. A bipolar signal is simply a signal that has positive and negative values. An audio waveform is an example of a bipolar signal and can in fact be used for the modulation signal to alter another audio waveform used as the carrier. The use of a bipolar signal causes an interesting difference in the output signal when compared to AM synthesis. When the modulation signal is used to quickly cycle the carrier signal's amplitude, like AM synthesis it forms two additional sidebands, but the carrier signal is cancelled out and disappears.

RM synthesis is used by [Vocoders](#) which are often used to effect a human voice's sound signal to create a "robotic" sounding variation.


Frequency Modulation (FM) Synthesis

Frequency Modulation (FM) Synthesis produces an output signal by oscillating the frequency of a source oscillator's signal. This process can generate fairly complex output containing multiple frequencies/sidebands with only two oscillators, requiring minimal computations. This computational efficiency is the reason for its invention and great popularity in earlier synthesizers and sound cards.

Wavetable Synthesis

Wavetable Synthesis is carried out by simply playing segments of digital audio to produce a realistic instrument or synthetic sound. The digital audio segments are stored as a table of waveforms in memory and played back at different speeds to produce

as a table of waveforms in memory and played back at different speeds to produce output of a different pitch for each musical note. One common addition to wavetable synthesis is that each instrument waveform contains a loop region. This region is started after the attack segment of the digital audio has played and repeats while the instrument's note is sustained (i.e. while a note is held). Then the release segment of digital audio finishes off the note. Wavetable synthesis in the form of envelopes and modulators are often combined with AM synthesis to add some variation to otherwise repetitive sounding output.



A wavetable instrument's waveform segments and example output.

Physical Modeling (PhM) Synthesis

Physical Modeling (PhM) takes a unique and complex, but more intuitive approach to synthesizing sounds. Synthesis is accomplished by simulating the physical properties of a real or fictitious musical instrument mathematically by defining exciters and resonators. Exciters are what trigger the physical model to start generating sound. Real examples include the hit of a drum, stroke of a bow, or pluck of a string. An input (such as a key press and the strength/velocity of the press) is translated into the appropriate parameters for simulating physical properties of instrument's input (such as blowing, amount of air, etc.). Resonators simulate the instrument's response to the exciter which usually define how the instrument's physical elements vibrate. Examples include, a plucked string, struck drum membrane, blown reed or even vibrating vocal chords. This type of sound generation can be extremely complex and require heavy computation, so many current available physical modeling synths use short-cuts or watered-down methods which enable them to respond in real-time.