12-2022

The Musical Implementation of Additive Synthesis

Quinn Ingrassia
California State University, Monterey Bay

Quinn Ingrassia

Dr. Lanier Sammons

MPA 475: Capstone

The Musical Implementation of Additive Synthesis

Sound synthesis can be executed using several approaches, each of which impart their own unique characteristics and pose their own unique limitations. Digital additive sine wave synthesis is one such technology that can perhaps produce the greatest number of sounds that are impossible to replicate by other methods. Modern music provides the best source for study of the aesthetic implementation of contemporary additive technology. This paper will outline some of the viable methods of musical implementation of additive synthesis by examining and deconstructing use cases in detail. But before we observe and deduce those methods of aesthetic implementation it will be useful to define additive synthesis generally, as well as the tools pertinent to replication of those methods.

Additive synthesis is simple in its most basic conception: it is any form of synthesis that adds together sine waves to create sounds. The system that organizes these sine waves could be called an additive synthesizer, a device that allows its user to add sine waves and adjust their frequencies and amplitudes.The Hammond organ, first patented in 1934, is an early example of an additive synthesizer (“The Hammond Story”). It uses a tonewheel to generate sine waves[1] and allows its user to control the output of the fundamental tone’s harmonics by adjusting the drawbars. As described in a blog post from Roland, who make products that emulate the original Hammond, “the original tone wheel organs use a series of mechanical tone wheels to generate sine waves that are the organ’s basic sound element. The tone wheels are

[1] Sine waves in the digital domain are mathematically perfect single frequency signals, however in the analog domain they are seldom so perfect. Sine waves become subject to noise and other distortions that smear the signal, making it difficult to replicate a pure sine wave in the real world. Many instruments that produce sine-like sounds produce more of a smeared sine wave.

physical spinning discs that rotate in front of something similar to a guitar pickup" (Watson). Sine waves are summed together in order to produce a sound like that of the pipe organ.

Contemporary additive synthesizers usually come in many different permutations and with many features but for the duration of this paper the focus will be placed on digital additive sine wave synthesizers. The paper will particularly focus on those which use sine banks: banks of numerous sine wave oscillators, holding dozens to thousands of individual oscillators. These synthesizers require many individual sine wave oscillators to recreate complex sounds by applying the logic of the Fast Fourier transform. The Fourier transform is a mathematical transform that can decompose the function of a complex wave into its component frequencies. The Fast Fourier Transform is an algorithm based on the original Fourier transform and more closely related to the discrete Fourier transform. Both of these editions of the Fourier transform process discrete sequences and decompose complex waves into sine waves, or they can be applied inversely (Inverse Discrete Fourier Transform or IDFT) to create a complex wave out of data ("Fast Fourier transform"). The fast Fourier transform is often preferred over the discrete Fourier transform when either might be appropriate because it is able to process data exponentially faster as the data size increases. Additive synthesizers use their sine wave oscillators as partials of the whole sound they intend to produce, functioning in a similar way to the IDFT. Additionally the most versatile additive synthesizers offer control of the frequencies, amplitudes, and phases of their individual sine wave oscillators. With those three parameters subject to manipulation, Pandora's Box is opened. The freedom this grants the user is deceptively vast. I will explain the importance of these features when discussing the synthesizers in detail. We will be studying two examples of synthesizers fulfilling the above criteria, and those are Image Line's Harmor and Native Instruments's Razor.

Harmor was released in 2011 and remains a popular choice for the enthusiasts of software additive synthesis. Many of the artists we will discuss later in the paper have documented publicly their use of Harmor, including Au5, L.M., and Voltra (Au5, "@Voltra Is the

Definitive Razor Sorceress Change My Mind"). Harmor mimics a subtractive synthesizer's interface. As Image-Line says, "Its modules will look familiar to subtractive synthesizer enthusiasts: oscillators, filters & phasers. These are featured in Harmor but, because they are performed on additive synthesis data, rather than audio, offer more freedom" ("Harmor"). The additive synthesis data mentioned here is what separates Harmor from the subtractive synthesizers from which it derives its interface. It refers to the data assigned to each sine wave oscillator's amplitude, frequency, and phase. I've already explained how the oscillators differ; in summary, an additive synthesizer uses only sine wave oscillators, which sum to produce the synthesizer's resultant sound. Harmor's filters and phaser modules alter the amplitude of these oscillators, simply boosting or attenuating the amplitude of the individual oscillators at designated frequencies in the complex waveform.

Standard filters attempt to apply the same effect, but what makes the additive version more liberating is a lack of artifacting. Many common subtractive filters like Image-Line's own Fruity Parametric EQ 2 use infinite impulse response (IIR) filters, which are recursive filters that usually impart a phase shift on the output signal ("Fruity Parametric EQ 2"). Although "it is easy to make an FIR (finite impulse response) filter have a linear phase [a filter type with functionally zero phase shift]… This is not the case with IIR (recursive) filters" (Smith 328). Because of this limitation, most filters will impact the phase of the signal. We can see below (fig. 1) the effect of Fruity Parametric EQ 2's minimum phase filter on a sawtooth wave playing the note C2 at 65.4hz:

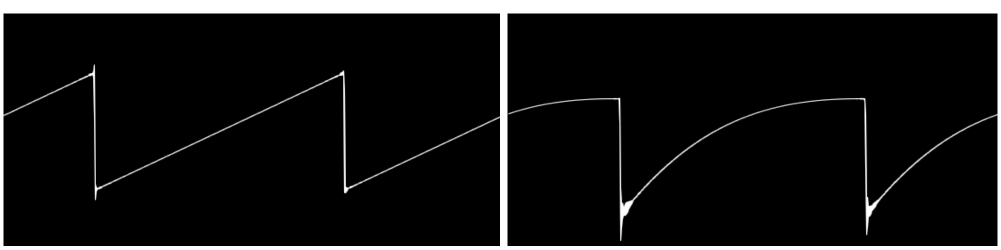

Fig. 1

The signal on both the left and right are the same, the only difference being Fruity Parametric EQ 2's high pass filter applied at 20hz to the signal on the right. Even though the filter cutoff frequency is well below the fundamental frequency of the wave, its phase is visibly altered. While the signal will sound the same to the human ear, this phenomenon manifests in more destructive ways in certain contexts.

One key example of those contexts is when a filter is swept through the frequency spectrum. The phase shift becomes audible as a pitch shift; after all, a change in phase of a signal is a change in the speed of the signal. Altering a signal's speed will cause a change in pitch, much like the Doppler effect.[2] If the user tried to apply a quick low pass filter envelope to their sound, the filter would alter the pitch as it sweeps through the frequency spectrum as seen below (fig. 2) in a spectrogram demonstration showing a plucked saw wave using a volume envelope (left) and a low pass filter envelope (right).

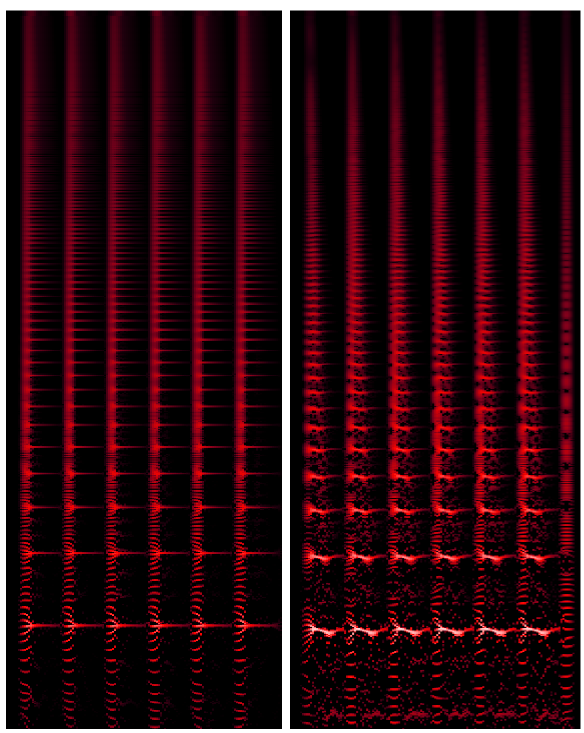

[2]"The Doppler effect or Doppler shift (or simply Doppler, when in context) is the change in frequency of a wave in relation to an observer who is moving relative to the wave source. It is named after the Austrian physicist Christian Doppler, who described the phenomenon in 1842" ("Doppler effect").

Fig. 2

The pitch alteration is visible here as a slight dip downwards. Since the filter is swept from high to low, the frequency of the signal at and around the cutoff is lowered. An additive synthesizer's filter is free of this limitation because the filter is essentially a curve or shape which tells the synthesizer where to boost, attenuate, or disable frequencies that cumulatively assemble the signal rather than adjusting the already constituted signal.

In fig. 3 I've captured Harmor playing a saw wave at the same frequency, but using its own filter.

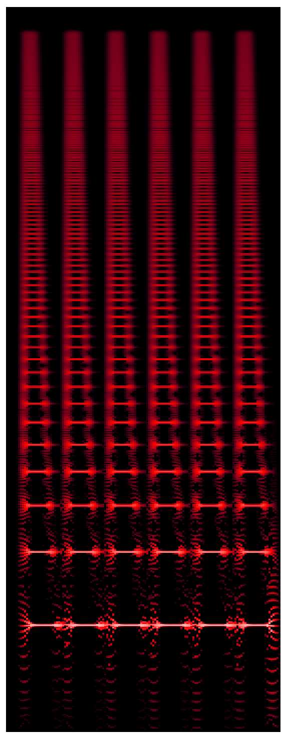

Fig. 3

There is no pitch change in this example. I chose to demonstrate the profound effect of phase shift with high and low pass filters because filters which cut frequencies out entirely also alter the phase most drastically. Additive filters are particularly powerful when it comes to cutting frequencies entirely because the side effect is not just minimized but completely eliminated. A low pass filter sweep, and any filter movement, using an additive filter can be applied as much and as fast as the user desires, with no consequence. Additive filters, then, are some of the most ideal in digital signal processing.

The additive filter is one application of the powers of the additive synthesis engine, but there are many more ways to make use of its unique qualities. Harmor has various components which take advantage of the additive engine including the harmonizer, blur, timbre module, harmonic phase module, frequency masking, filter resonances, pluck module, unison, equalizers, detune module, and more. But for now it's most pertinent to discuss the phaser module and save the rest for when we discuss applications in music. The Harmor developers look at phasing as "the process of creating constantly moving frequency cancellation/s in a sound" ("Harmor"). This is indeed the classic effect of a phaser, which mixes a copied signal with an all-pass filter applied to it back into the original signal. The all-pass filter rotates the phase of the copied signal and when mixed back with the original signal it creates a series of constructive peaks and destructive notches in the frequency spectrum. Harmor's phaser skips a few steps and simply applies a pattern of notches according to preset envelopes or even a user defined phaser envelope. These phaser envelopes are the shapes of the notches which will be cut out of the waveform. Below (fig. 4) are three examples of phasers, from left to right: the classic phaser, Harmor's phaser, and Harmor's deep phaser type which showcases the profound notching that is possible.

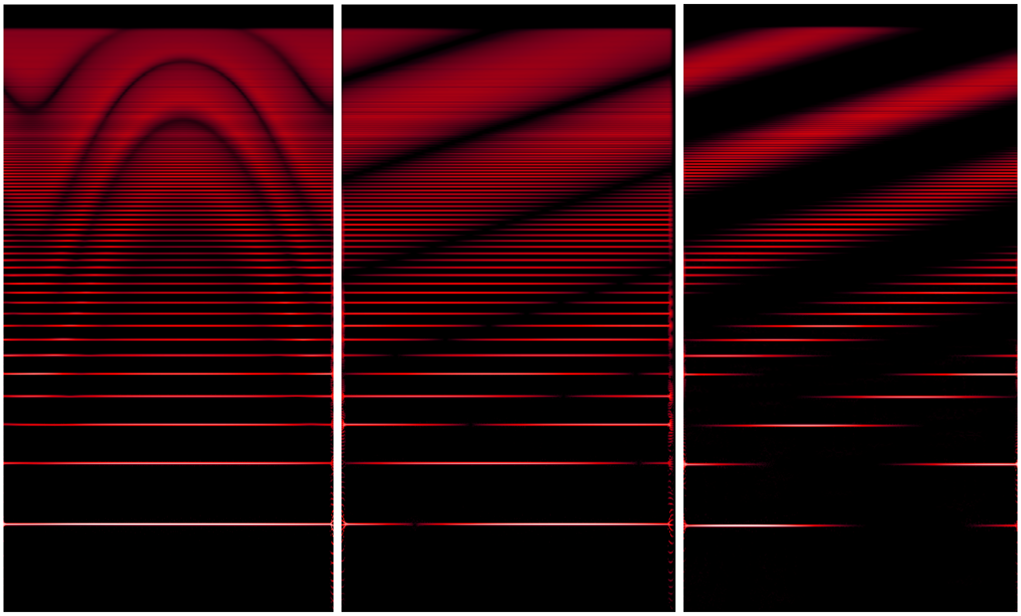

Fig. 4

Harmor is able to cut frequencies more deeply than a traditional phaser because its phaser works like its filter, attenuating, boosting, or disabling frequencies in the waveform according to what the user desires. Harmor's phaser also allows precise control of the movement speed, movement direction, offset of the phaser position, width, mix or depth, and scale to octave, Hertz, or harmonics. A change to any one of these parameters can alter the output significantly, and in my experience the phaser is the quickest path to fascinating sounds.

Razor is similar to Harmor in that it is an additive synthesizer released in 2011, but it offers a plethora of its own valuable features. Razor may be the most relevant utility in the additive synthesis landscape because it is widely available as part of Native Instruments Reaktor in any digital audio workstation or standalone. Razor offers four primary additive features: its oscillators, filters, dissonance effects, and stereo effects (Errorsmith). Each of these are unique to Razor and will constitute the basis of its sounds. The oscillators have fourteen available waveforms that can be warped according to set parameters. The most versatile is the "pulse to saw" waveform, which will sound most like classical subtractive synthesis, however the other options are well worth using. For example the primes waveform is uniquely additive, playing only prime number harmonics, which would be difficult to achieve subtractively. Two filters can be applied in sequence to modify the oscillators with a choice of 18 filter types. They function very similarly to Harmor's, boosting, attenuating, or disabling harmonics at designated frequencies. The additive effects are a special feature of Razor. While some of the effects like beating or autopan are meant to substitute for simple post processing effects found elsewhere, others like centroid and resonant reverb can profoundly alter the sound, and each exists only in the Razor synthesizer.

Before I get into musical implementation of additive synthesis, I should address a meaningful limitation to additive synthesis. In figures 3 and 4 there is a visible frequency cap at around 17khz for a note played with its fundamental at C4, and a cap around this area is persistently common among additive synthesizers because the partial count reaches its limit here. Since each additive synthesizer has a certain number of partials in its sine bank, it can only replicate that number of harmonics in the complex waveform it means to construct. Harmor has 516 partials ("Harmor"), and Razor has 320 (Errorsmith). These frequency caps are often high enough that they present no real problem to the user and would be difficult to detect anyway. However, the partial cap drops progressively until it hits a minimum of 16Khz at C0. A bass note at C0 contains more spectral content between the fundamental and the frequency cap than a note at C4. The limited partial count means that the fidelity of the signal decreases as the synthesizer runs out of partials to allocate to harmonics. While this may not sound like a significant difference, the effect on the sound seems to result mostly from the ratio of harmonics to the fundamental frequency according to my own analysis.

Harmor is locked at 516 partials maximum, however Razor has its inner workings exposed in its Reaktor ensemble. From there the user can manipulate the partial count; the most concise method by which to do so can be found here from a tutorial video by Au5 (Au5, "Unlimited Partials: Razor Hack"). Pushing the partial value to the thousands can repair those lossy bass sounds and impart a glossy quality to the sounds since most of the additional fidelity manifests in the high frequencies. I will recommend a maximum of 2560 partials. Partial counts are limited in the first place because processing hundreds to thousands of oscillators can drastically impact the user's CPU performance. Processing power in 2011 when Harmor and Razor were released was not adequate to reproduce much more than 500 sine waves. Modern computers should be able to handle more though, and 2560 partials (8x Razor default) should suffice for any use case since it caps out over 20khz at C0.

The sounds of additive synthesis can contribute in many interesting ways to music. They have certainly found a home in electronic music as some of the previously mentioned artists use additive synthesis adeptly and showcase its unique sounds. While additive synthesis is able to reproduce sounds similar to those popularized by other more common types of synthesis, I term the sounds most easily and intuitively accessed through additive synthesis "additive native sounds," and these are the subject of this paper's remaining study. I've selected a few categories of sounds named by myself, or by their author if the documentation exists, to be the focus of excerpts aimed at determining their musical purpose derived from the context of the songs in which they are used, as well as a method to recreate them in Harmor or Razor.

1. Bass Plucks:

Because of the clean filtering made possible by additive filters, bass plucks are made simple. In his Always In A Nightmare breakdown, Au5 says "this is probably the cleanest supersaw pluck bass that you can make," (Daw Nation, "Making Emotional Bass Music"), and he discusses his use of additive synthesizers for bass plucks in his Black and Yellow remix breakdown (Daw Nation, "Making a Dubstep Remix of Black and Yellow: AU5") and his Potions remix breakdown (Daw Nation, "Slander & Said the Sky - Potions (Au5 Remix) Full Walkthrough"). Bass plucks tend to be excellent rhythmic backbones to underlie a harmonic progression. Reducing pitch distortion is ideal for a harmonically rich sound of this type because there are fewer frequencies in the low frequency range due to the logarithmic nature of human frequency perception. Each individual frequency is relatively more important than in higher frequency ranges because of this. For example, unison detune decreases pitch fidelity to the human ear less as the note becomes a higher frequency because it perceives a smaller gap between each harmonic, therefore a detune value of a certain number of cents becomes less significant. Since phase alterations from filter sweeping can cause pitch distortion in the low end, the artist may want to avoid this by implementing an additive-based solution. Harmor and Razor can both modulate their filter cutoffs with envelopes. Simply applying this envelope to some sort of low pass or band pass filter shape will do the trick.

2. Au5's Fields of Gray Bass:

Au5's Remix of Fields of Gray by Infected Mushroom (Au5, "Infected Mushroom - Fields of Grey (AU5 Remix)") was the earliest instance I could find of this sort of bass heard at 2:01. It has a unique character, sounding almost liquid because of the phase alteration one might find when applying an all pass filter to a harmonically rich wave, but also metallic because of the atonal harmonics present. This sort of sound seems not to occur in nature, so using it in music could evoke a futuristic or alien aesthetic. While there is seldom only one way to produce a certain type of sound, the simplest method for this bass is through Harmor, as far as I know. Starting from Harmor's default preset, only a few changes need to be made, and the default unit order will work fine in this instance. The phaser will need to be turned on with its scale mode set to harmonics. While the octave and Hertz modes tell the phaser to space notches evenly per octave or Hertz respectively, the harmonics mode spaces oscillator cancellation(s) based on the harmonic's individual frequency. By applying one of Harmor's units to alter harmonic frequency, the phaser will respond. A simple way to achieve this effect is by using Harmor's prism unit, adjusting the harmonic prism shape, the prism's parameters, and the phaser's parameters to find the desired bass quality. This process can yield many sorts of sounds, but it's not difficult to find the kind of liquid timbre of Au5's "Fields of Gray" bass. The metallic timbre can be produced using Harmor's harmonizer unit. Randomizing the shift and gap matrix and adjusting the mix to taste will complete the sound. If the parameters in each step are adjusted carefully to emulate the "Fields of Gray" bass, the resultant sound should be similar or even nearly identical. Playing with this method can also yield profoundly different results, any number of which could also be musically useful.

3. L.M. Permafrost Bell:

In L.M.'s track "Permafrost," the drop contains a bell-like bass sound. This bass seems to be a layer of both an additive bell sound and what is perhaps some sort of frequency modulation (FM) bass layer. Together they are being distorted, and the low frequencies are amplitude modulating the higher frequencies. The final bass sound is unlike anything else I've heard during my research and while the FM layer is not something I have yet figured out, I can offer insight into the additive bell layer. Additive synthesizers tend to excel in creating bell-like harmonics because they can usually randomize their own harmonics. Bell sounds are full of overtones which have no harmonic relationship to each other, so randomizing harmonics gets the job done. I suspect this particular "Permafrost" sound was made in Razor where a simple 3 step process will create a sound somewhere close. There is a pitch envelope on the "Permafrost" bass so applying a wide sweep downward of 60 semitones (the maximum) imparts the same effect. The "Pulse to Saw" oscillator will work here, set to saw. Optionally, applying an envelope to the pulse/saw parameter can enhance the sound. Lastly, the semitone spacing dissonance effect is the crux of the sound and brings it up to the original's signature by rearranging the harmonics. By applying an envelope to the spacing amount and adjusting the amount and pitch parameter positions to find the appropriate placement, the sound is complete. The "Permafrost" bell dominates the soundscape when it is introduced. As the semitone spacing amount is increased, it changes the position of the partials in the frequency spectrum, but not the amplitude. This means that if a saw wave is spaced out so that its first 20 harmonics occupy the same range in frequency as the entire saw wave (below Nyquist) did before, the higher frequency harmonics will be much louder since their original amplitudes have not been altered. This seems to make the bell feel louder than a typical sound playing the same note. This sound is useful as an oppressive presence in music.

4. Voltra Showcase Pluck:

Voltra seems to be an adept user of digital additive sine wave synthesizers. She incorporates unique sound design into her music, and this has become part of the Voltra brand. Beyond what my ears tell me when I listen to her songs, upon analyzing them using a spectrogram, it seems that some of her sounds are "additive native," and indeed the simplest way to reproduce them is using additive synthesizers. In her "2022 Showcase" (Voltra), Voltra uses a pluck at 5:35 that caught my ear (fig. 5, left).

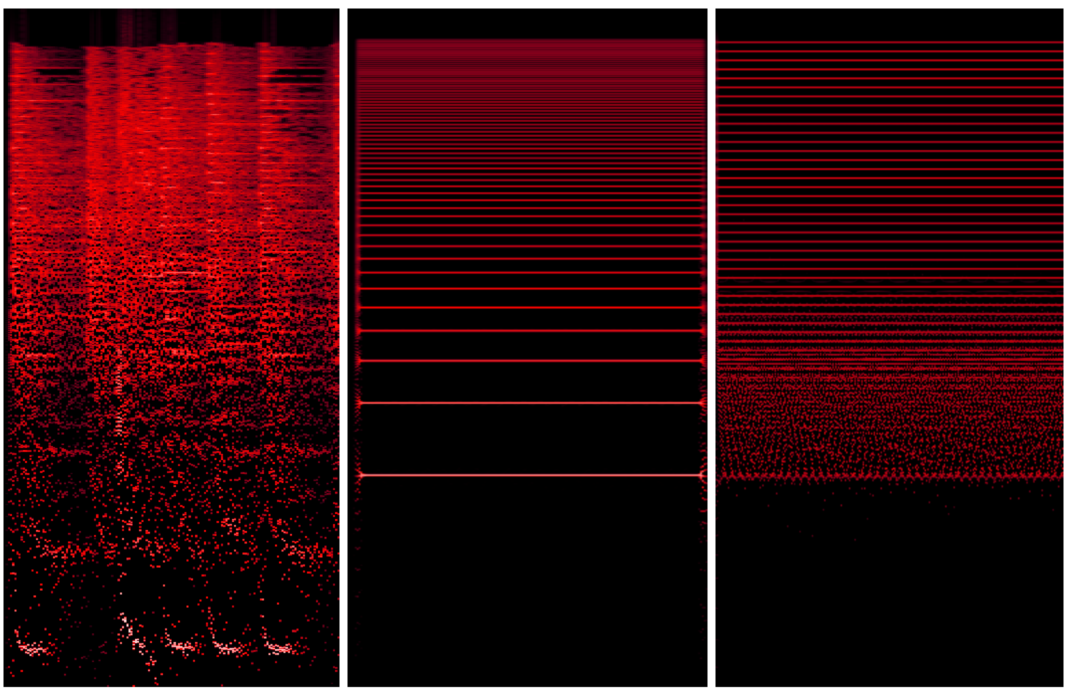

Fig. 5

In this image the pluck is what appears to be evenly spaced harmonics from the mid to high frequencies. Evenly spaced harmonics like this are rare among usual wave shapes. A saw wave for example (fig. 5, center) becomes denser as its harmonics reach higher frequencies when the spectrogram is weighted to space octaves evenly. Essentially, a saw is spaced evenly to Hertz weighting and the pluck in question is spaced evenly to octave weighting. Harmor is capable of easily recreating the even spacing found in Voltra's showcase pluck by engaging its octave pitch distribution (fig. 5, right). The original pluck has a soft pitchy quality, which can be emulated by applying a volume envelope for the overall shape and a pitch envelope to mimic the pitchiness. The "Voltra 2022 Showcase" pluck seems to have gaps at certain frequencies at random times. This can be done using the phaser unit or even the unison unit to create the appropriate frequency cancellations. There is one last attribute to this pluck that I suspect to be a layer: the sub. To recreate this sound faithfully it is simplest to layer the sub bass separately, with its own pitch envelope, then route each sound to one bus and distort them. This sound, again, evokes an unnatural feeling since the same sort of sound rarely exists naturally. The "Voltra 2022 Showcase" pluck in particular is aggressive and sharp. Since the harmonics of this sound are spaced so consistently up to the very highest frequencies, it will not face frequencies masking problems as often as will other instruments in the mix. It also survives distortion well, maintaining its character even after a significant amount is applied to it. These qualities make the "Voltra 2022 Showcase" pluck a solid choice for clean stabs that convey a sense of precise force.

The four examples deconstructed above are a few of the infinite sounds additive synthesis makes possible. While additive synthesis is not free of limitations, it continues to hold great potential as technology improves. There are certainly more features to implement or simply to combine from each synthesizer to create an extremely versatile new one. The goal of this paper was to discover the uses and possibilities made available by additive synthesis and introduce those potentials to the readers. While it's useful to learn more about one type of synthesis for its own sake, learning about one type can also influence and improve your understanding of what is and isn't possible with other types. Additive synthesis offers the most molecular perspective on the creation of sound and is a powerful tool both to learn about synthesis and create with synthesis. There is no sound completely inappropriate for music: anything can be sonified. But there are those sounds which are closer to what culture has imprinted onto the collective schema as "instruments," and those which are further. "Additive native" sounds tread the line where something becomes a sound or an "instrument," while additive synthesis generally is also excellent at recreating instrument-type sounds. This unique quality is exactly what makes additive synthesis effective in music.

Works Cited

Au5. "Infected Mushroom - Fields of Grey (AU5 Remix)." *SoundCloud*, https://soundcloud.com/au5/infected-mushroom-fields-of-grey-au5-remix?si=67513c1b00bf40b8aa28b20409d63a15&utm_source=clipboard&utm_medium=text&utm_campaign=social_sharing.

Au5. "Unlimited Partials: Razor Hack." *YouTube*, 30 Nov. 2019, https://youtu.be/xvVedjkvFgM.

Au5. "@Voltra Is the Definitive Razor Sorceress Change My Mind." Twitter, 11 May 2021, https://twitter.com/Au5music/status/1392243533780688899?s=20&t=EHGrbJA9-JpqKDzIg-L1wQ.

Daw Nation. "Making a Dubstep Remix of Black and Yellow: AU5." *YouTube*, 29 Apr. 2021, https://youtu.be/oLBqmi0ot_g?t=260.

Daw Nation. "Making Emotional Bass Music: Au5 in the DAW: Always in a Nightmare." *YouTube*, 1 May 2020, https://youtu.be/aaImnBvznRU?t=1787.

Daw Nation. "Slander & Said the Sky - Potions (au5 Remix) Full Walkthrough." *YouTube*, 16 Apr. 2020, https://youtu.be/bGosyV43yDQ?t=619.

"Doppler effect." *Wikipedia: The Free Encyclopedia*. Wikipedia, The Free Encyclopedia, 22 Nov. 2022. Web. 8 Dec. 2022.

Errorsmith. "Native Instruments Razor." *Native Instruments*, March 2011, https://www.manualslib.com/manual/540823/Native-Instruments-Razor.html?page=74. Accessed 8 December 2022.

"Fast Fourier transform." *Wikipedia: The Free Encyclopedia*. Wikipedia, The Free Encyclopedia, 19 Nov. 2022. Web. 26 Nov. 2022.

"Fourier transform." *Wikipedia: The Free Encyclopedia*. Wikipedia, The Free Encyclopedia, 24 Nov. 2022. Web. 26 Nov. 2022.

"Fruity Parametric EQ 2." *FL Studio 20 Reference Manual*, Image-Line Software, 13 Jan. 2021, https://www.image-line.com/fl-studio-learning/fl-studio-online-manual/html/plugins/Fruity%20Parametric%20EQ%202.htm.

"The Hammond Story." *HammondOrganCo.com*, https://hammondorganco.com/the-hammond-story.

"Harmor." *FL Studio 20 Reference Manual*, Image-Line Software, Oct. 2011, https://www.image-line.com/fl-studio-learning/fl-studio-online-manual/html/plugins/Harmor.htm.

L.M. "Permafrost." *SoundCloud*, https://soundcloud.com/rushdownrecs/lm-permafrost?si=90f28fabd31b4f2590f6ed70653c516c&utm_source=clipboard&utm_medium=text&utm_campaign=social_sharing.

Smith, Steven W. *The Scientist and Engineer's Guide to Digital Signal Processing*. California Technical Pub., 1997.0

Watson, Mark. "What Is Virtual Tone Wheel Technology?" *Roland Resource Centre*, Roland, 3 July 2020, https://rolandcorp.com.au/blog/what-is-virtual-tone-wheel-technology#:~:text=The%20original%20tone%20wheel%20organs,of%20a%20vibrating%20tuning%20fork.

Voltra. "Voltra 2020 Showcase." *SoundCloud*, https://soundcloud.com/voltra/2020-showcase?si=4128121d8ccc468f8d95d62bd8027f09&utm_source=clipboard&utm_medium=text&utm_campaign=social_sharing.