Design, Generation, and Analysis of Selected Musical Sounds Using a Modular Synthesizer

Pre-scientific paper by

Faveo Hörold

class 8AR

supervisor: Kathrin Glantschnig

March 4, 2018

BG|BRG Villach St. Martin

9500 Villach, St.-Martiner-Straße 7

Abstract

The purpose of the work presented here is to ascertain how well the sounds of acoustic instruments can be recreated with an analog modular synthesizer. To answer that question, several synthesizer patches were recorded and compared to recordings of real instruments. A patch is a configuration of an analog synthesizer that produces a specific sound. The sounds resulting from these patches were analyzed using both objective and subjective methods. It turns out that some families of instruments are easier to synthesize than others, depending on the amount of physical factors that interact with each other in the real instrument to produce the sound. Instruments with less interacting elements are far easier to synthesize than instruments with several interacting parts. Another conclusion that can be drawn from the results is that very good resemblances of real instrument sounds can be created with exceptionally simple methods. Additionally, proof is presented that many synthesizer patches can be generalized and used to replicate different instruments with similar properties. For example, the patch for a flute sound can easily be tweaked to produce an oboe or clarinet sound.

Contents

1 Introduction

Analog music synthesizers are experiencing a comeback due to new, cheaper production methods and superior performance to digital synthesis methods in many areas. Analog synthesizers do not need complex software or expensive soundcards for real time performance, and digital software can only so much as approximate truly analog sounds. My father and I have begun designing and constructing a do–it–yourself (DIY) analog modular synthesizer. Such an instrument consists of many modules which each take an input signal and modify it in a very specific way to produce the desired output signal. A major challenge in analog synthesis is finding ways to interconnect modules and adjust their settings so that they produce a specific sound. Such a group of interconnected modules for creating a specific sound is called a patch. The modular analog synthesizer used for the work can be seen in figure 1. A patch has been set up to demonstrate the interconnectivity of the modules.

There are two main questions that I aim to cover in this work. First: "How is a synthesizer patch that reproduces a real world sound created?" Second: "What similarities and differences are there between real and synthesized sounds?" Of course, it is not possible to create accurate emulations of musical instruments with an analog synthesizer, but approximations that sound similar are feasible.

Furthermore, I will present several patches for generic acoustic instruments and subsequently analyze the patches themselves and the sounds they produce. The patches are from several different sources and will be compared both in terms of functionality and audible results.

The goal of this work is to find out how well analog synthesizers can be used to recreate actual instruments. Another important objective is to prove that a single synthesizer patch can be used for a group or even family of similar instruments.

First, several concepts of synthesis will be presented and explained, and several instruments will be selected for examination. After this, the main part of this work will deal with describing, implementing, and analyzing the various patches for each instrument. Then, I will summarize the results of

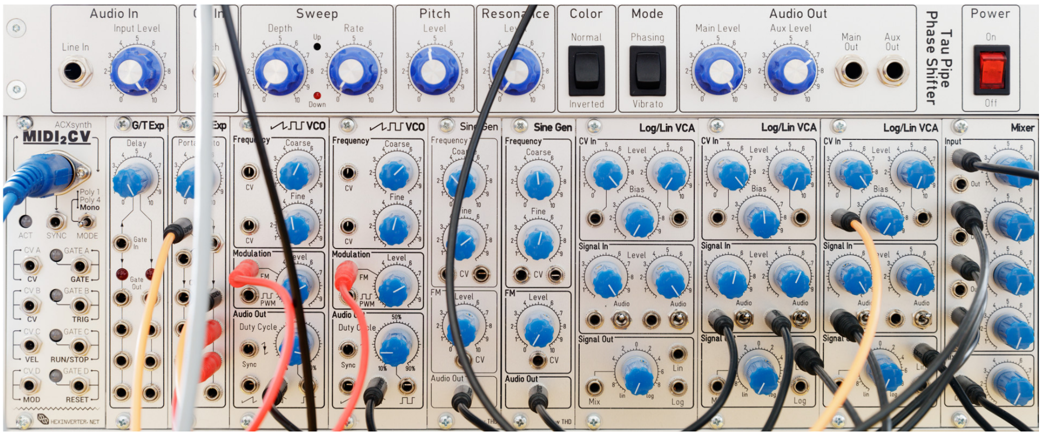

Figure 1: A closeup photograph of the analog modular synthesizer used for producing the various sounds presented in this work.

the experiments and assess the complexity and realism of the various patches in context.

2 Background Information and Preparation

2.1 Classification of Instruments

The aim of this work is to create synthesizer patches that are not specific to one single instrument but can be easily tweaked and adapted to be used for a wide range of similar instruments. This means that a single patch must be able to generate the sounds of several different instruments without changing the underlying structure of the patch. Thus, it must be decided which instruments are similar enough to be reproducible with a single synthesizer patch. The most commonly used method to classify instruments today is the Hornbostel–Sachs method which groups instruments by their method of sound generation (Ziegenrücker, 2009, p. 243). The creation of most synthesizer patches in this work relies on the analysis of the recreated instrument to some extent. Most patches also simulate the physical behavior of an instrument. For this reason, instruments that are grouped together in the Hornbostel–Sachs classification system will also have similar synthesizer patches, and in many cases, one patch can be used to synthesize several instruments. At least one instrument per major category in the Hornbostel–Sachs classification method has been synthesized. The category of idiophones is large and diverse, so I have chosen to synthesize the xylophone because it is one of the most common representatives of its family. The instruments are presented in the following list.

- Aerophones: Orchestral flute
- Idiophones: Xylophone
- Membranophones: Bass Drum
- Chordophones: Acoustic guitar (plucked), Violin (bowed)

2.2 Introduction to Synthesis

2.2.1 Definition of Sound

Sound is any physical oscillation audible to our ears (Paul A. Tipler, 2004, p. 465). Such an oscillation can be described in many ways. Most commonly, the amplitude of the sound's waveform is plotted against a time axis. This is what is usually displayed by an oscilloscope when an analog audio signal is connected to the input and is called a waveform. Another method for displaying the information contained in a sound is to plot the amplitude of all frequencies that are found in a section of the sound using fourier analysis (Ekbert Hering, 1989, p. 374). In this work the Fast Fourier Transform (FFT) algorithm is used to generate all of the fourier analysis diagrams, which will also be referred to as "FFT" diagrams for convenience.

An FFT can describe any oscillation because it defines the exact amplitudes of all frequencies contained in the oscillation at a given time. A waveform can also describe any oscillation, but an FFT is usually much more useful as it displays frequency information in an easily readable format. For audio purposes, an FFT must only define the amplitudes of frequencies in the human auditory range, which, according to Tipler (2004), ranges from about 20 Hz to 20 kHz (Paul A. Tipler, 2004, p. 493). All of the FFT diagrams in this work will cover a smaller frequency range than this because frequencies near the limits of human hearing are not of particular interest when trying to recreate the sounds of common musical instruments with analog synthesizers.

2.2.2 Analog Synthesis

There are two main methods by which synthesizers produce sound. The less common method, called additive synthesis, involves the addition of fundamental frequencies (sine waves) to construct a sound (Welsh, 2006, p. 4). The more common method, called subtractive synthesis, involves the subtraction of unwanted frequencies from an already rich spectrum (Welsh, 2006, p. 5). Much like a sculptor removes unwanted chunks from a stone block, the synthesist removes unwanted ranges and chunks of frequencies from a frequency spectrum. Just like the hammer and chisel is the tool of the sculptor, the filter is the tool of the subtractive synthesist.

While these are the basic concepts, many additional tricks are applied to alter sounds. Many of these cause variations in frequency and amplitude over time. For example, different frequencies can appear and disappear, as well as change in amplitude, making the spectrum dynamic instead of static.

The synthesizer used in this work is controlled by a keyboard. Several control signals are used to regulate how long and loudly which note is being played on the synthesizer. These signals are called pitch control voltage (pitch CV), gate, and velocity. In this work, the term control voltage (CV) by itself refers simply to pitch CV, even though gate and velocity are also CVs. According to Vail (2014), pitch CV is most commonly calibrated to a one volt per octave response (Vail, 2014, p. 139). Thus, if the note A4 (440 Hz) corresponds to a pitch CV of 4 V, then A5 (880 Hz) will correspond to 5 V. Another signal, called gate, is used to control whether a note is on or off. Gate is a binary signal, which means that it can be either high or low. According to Vail (2014), a high signal means that a note is being played and a low signal means that no note is being played (Vail, 2014, p. 130). A third signal, called velocity, specifies how loud a note is being played (Vail, 2014, p. 12). A larger voltage indicates a louder note.

In the setup used for this work, the Musical Instrument Digital Interface

(MIDI) (Vail, 2014, p. 16) signal from a digital keyboard was fed into a MIDI to CV converter which outputs pitch CV, gate, and velocity as analog signals. Although a digital keyboard can play many notes simultaneously, this is not usually the case for an analog synthesizer. Most analog synthesizers are monophonic or duophonic, meaning that they can only produce one or two notes simultaneously. One of each control voltage is needed for every extra note, leading to large and complex setups, which is why polyphonic analog synthesizers are quite rare. In the case of this work, only one note was desired for each sound sample, so a monophonic setup was sufficient.

The voltage controlled oscillator (VCO) is usually the starting point for making any sound with a synthesizer. The oscillator's pitch is controlled directly by pitch CV. The pitch knob on a VCO allows the synthesist to add an offset to the pitch CV, thus making it possible to tune the oscillator to a desired pitch. Welsh (2006) explains that VCOs usually supply simple waveforms such as sawtooth, pulse and triangle (Welsh, 2006, p. 5). Sine wave VCOs are also sometimes used (Vail, 2014, p. 137). A sawtooth wave's frequency spectrum is a harmonic series of overtones. This means that if a 440 Hz sawtooth wave is generated, it will be a mixture of sine waves at integer multiples of 440 Hz. A corresponding FFT plot will contain a series of regular peaks starting at 440 Hz. As can be seen in figure 3, the frequency spectrum of a sawtooth wave does not change if the waveform is inverted. A pulse wave's spectrum is just like that of a sawtooth, but with harmonics only at odd multiples of the fundamental frequency. The waveform of a pulse wave repetetively jumps between a low and high state. For this reason, such a wave is often also called a square wave. According to Welsh (2006), the duration that the signal of the waveform stays in one state opposed to the other is called duty cycle or pulse width, and a wave with equally long high and low parts is said to have a pulse width of 50% (Welsh, 2006, p. 7). When the pulse width is changed, the length of a section changes, which causes the amplitudes of the peaks in the spectrum to change. New peaks will appear at odd multiples of the fundamental frequency when the pulse width is at any value other than 50%. A wave with a pulse width of 0% or 100% is silent because it is constantly either high or low. A triangle wave's spectrum is like that of a pulse wave with a 50% pulse width, but the higher–pitched harmonics' amplitude decreases more quickly. The frequency spectrum of a perfect sine wave consists of a single frequency peak at the sine wave's frequency. All of these waveforms have been plotted in figure 2. Corresponding FFT plots are in figure 3. For completeness, white noise was included in these figures. White noise is usually produced by a seperate module and not by a VCO because it is not an oscillation. Welsh (2006) demonstrates that an FFT plot of white noise shows a uniform distribution of amplitudes across the entire spectrum of frequencies (Welsh, 2006, p. 12). As Welsh (2006) explains, another feature of a VCO is that it can be synchronized to another oscillator, called the master oscillator. The synchronized oscillator is called the slave oscillator. Whenever the master oscillator completes one cycle, the cycle of the slave oscillator is reset. Synchronization creates dips and valleys in the frequency spectrum of the slave oscillator (Welsh, 2006, p. 9).

The low frequency oscillator (LFO) is simply a VCO operating in a very low, usually subaudible, frequency range (Vail, 2014, p. 138). Some oscillators can be used both as LFOs and VCOs if their frequency range is large enough. LFOs are usually used as generators for control voltages that are further used for generating slow vibratos or similar effects. The waveforms of the LFO used in this work are pulse, sawtooth, inverted sawtooth, and triangle.

The voltage controlled filter (VCF) is arguably one of the most important modules in subtractive synthesis because it enables the precise manipulation of an existing sound. According to Vail (2014), VCF's main purpose is to attenuate parts of an input frequency spectrum (Vail, 2014, p. 21). The main filter types are lowpass, highpass, and bandpass. In a lowpass filter, high frequencies are attenuated while low frequencies are passed unchanged. In a highpass filter, low frequencies are attenuated while high frequencies are passed unchanged. A bandpass filter is essentially the combination of a lowpass and highpass filter in the sense that frequencies above and below an unchanged frequency band are attenuated. There are also some less used and more exotic filters, such as the notch filter (also called band–reject filter (Welsh, 2006, p. 17)), which is the exact opposite of a bandpass filter,

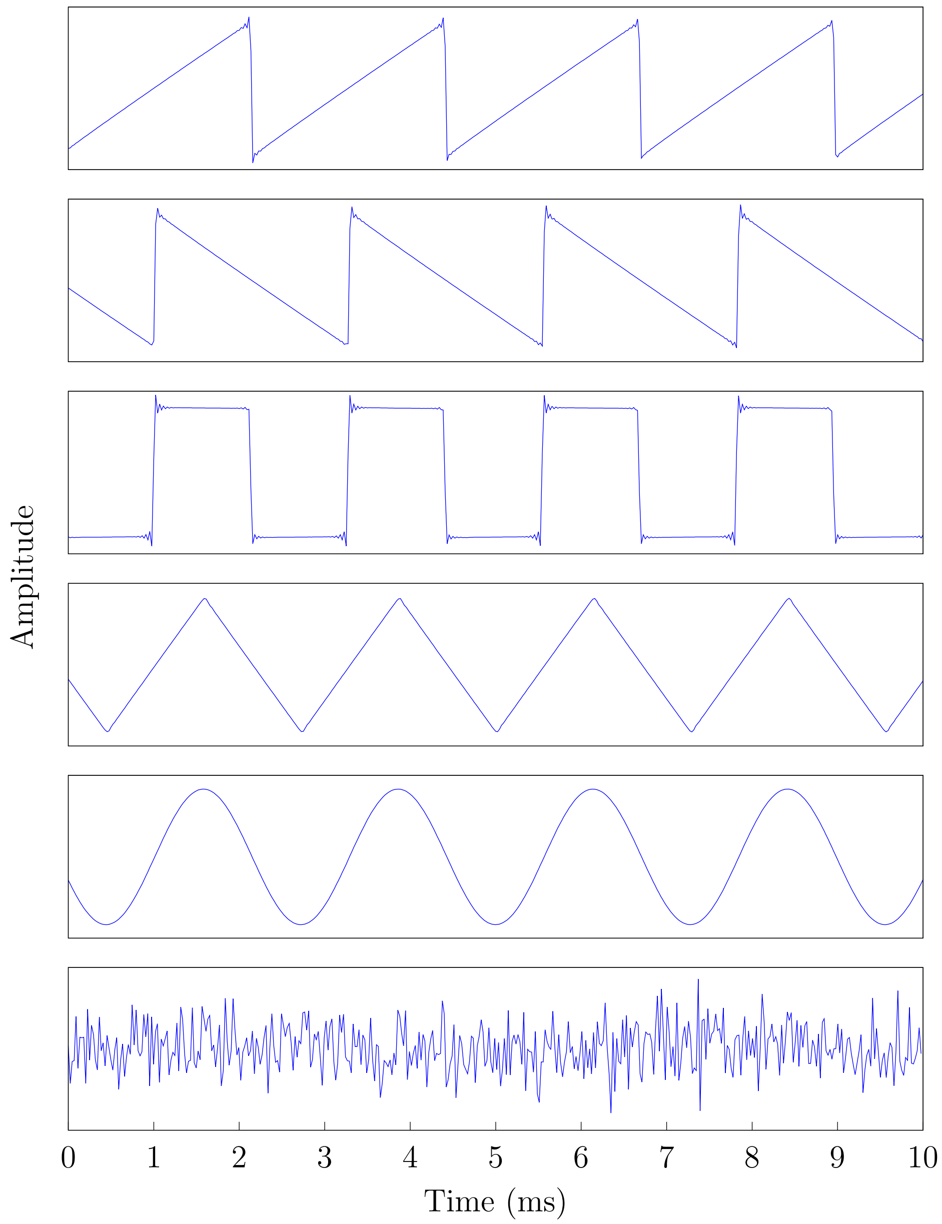

Figure 2: A collection of plots of various basic waveforms. From top to bottom these are: ascending sawtooth wave, descending sawtooth wave, pulse wave, triangle wave, sine wave, and white noise. See figure 3 for the corresponding FFT plots.

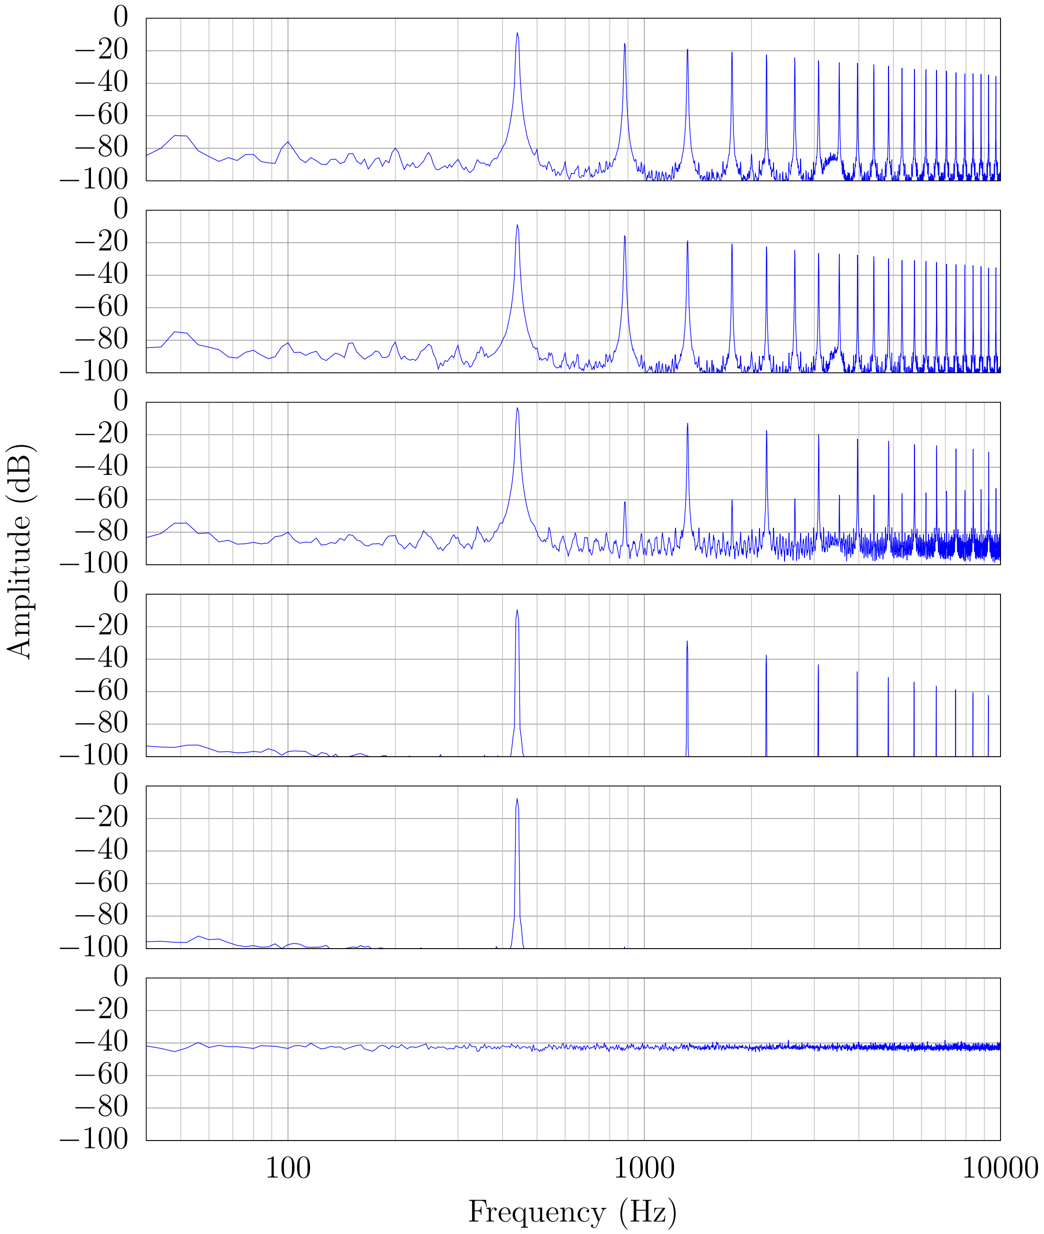

Figure 3: A collection of FFT plots of various basic waveforms. From top to bottom these are: ascending sawtooth wave, descending sawtooth wave, pulse wave, triangle wave, sine wave, and white noise. See figure 2 for the corresponding waveforms. Note that the ascending and descending sawtooth waves have identical spectra.

attenuating only a range of frequencies and letting everything else pass. As described by Welsh (2006), the frequency at which filtering takes place is called the cutoff frequency (Welsh, 2006, p. 17). For low- and highpass filters, this sets the divide between low and high frequencies. For band and notch filters, it sets the frequency at which the filter should act. It is physically impossible to perfectly cut off a specific range of frequencies. Thus, frequencies are gradually attenuated more and more as they move farther away from the cutoff frequency. According to Welsh (2006), the term filter slope describes how well frequencies are attenuated (Welsh, 2006, p. 19). Another feature of filters is resonance. In lowpass and highpass filters, resonance amplifies frequencies near the cutoff point. In bandpass and notch filters, the resonance setting changes the width of the cutoff band (Welsh, 2006, p.19).

The voltage controlled amplifier or attenuator (VCA) changes the amplitude of an input signal depending on a CV signal (Vail, 2014, p. 151). VCAs are essential for almost every single synthesizer patch, as they turn a note on and off when a key is pressed. Every patch presented in this work has a VCA at the end of the signal chain to control the amplitude of the output sound.

An envelope generator creates an envelope CV when it receives a high gate signal. A simple attack–decay (AD) or attack–release (AR) envelope generator ramps up the output signal when the gate CV is high, and ramps it down when the gate CV is low. The attack and release parameters define how long each respective operation takes. The attack time defines how long the CV takes to ramp up, while the release time defines how long the CV takes to ramp down. The more complex attack–decay–sustain–release envelope generator (ADSR) (Vail, 2014, p. 152) adds the two extra parameters decay and sustain. Sustain is the level at which the output envelope CV should stay while gate CV is high. In the attack phase, the signal always ramps up to 100% . Thus, the decay setting is needed to define how long it takes for the output CV to go from 100% to the possibly smaller sustain level. According to Welsh (2006), the signal generated by an envelope generator can be used to control many things, but the most important use is amplitude control. By sending an envelope CV into a VCA, the loudness contour of a note can be controlled over time (Welsh, 2006, p. 22).

According to Welsh (2006), a ring modulator multiplies two input signals. When these input signals are two sine waves, the output is the sum and the difference between the frequencies of those sine waves (Welsh, 2006, p. 15). It is very difficult to create harmonic sounds with a ring modulator, as every frequency in each input spectrum is added and subtracted with every other frequency. Instead, metallic bell–like sounds are more common.

A fixed filter bank contains several bandpass filters with fixed cutoff frequencies (Vail, 2014, p. 18). Large fixed filter banks are quite uncommon, so this module was avoided and only used in the patch for the violin. Fixed filter banks allow precise manipulation of a frequency spectrum, which is why they are very good for simulating physical resonators with special properties, like the body of a violin or guitar. The fixed filter bank used for this work contains eighteen bandpass filters with a slope of 24 dB each. The filter frequencies are evenly spaced in a range from 120 Hz to 7 kHz, as can be seen in figure 5.

2.3 Description of Analysis Procedure and Resources

As described in 2.2.1, sound can be fully described by an FFT spectrum. This is one of three methods that will be used to compare the synthesized sounds of each patch with the recordings of real instruments. The FFT will be an average of the complete sound as it changes over time, thus losing any information about changes that happen to the spectrum as the sound evolves. However, none of the synthesized patches have massively varying frequency spectrums, so it has been assumed that any information losses stemming from the use of an averaged FFT are minimal. However, information on the loudness of the sound over time is lost completely, so the amplitude envelope of the sound will be used as a supplemental metric. According to Vail (2014), an amplitude envelope is simply a function that shows the loudness of a sound over time (Vail, 2014, p. 20). The last method to be used is a subjective one. I will compare recordings of the synthesizer patch with recordings of the real instruments and attempt to describe differences and similarities as I hear them. In summary, it will be assumed that each sound has a static frequency spectrum whose general amplitude changes over time. Thus, a single FFT and amplitude envelope will be used to objectively describe a sound. Any audible changes in the frequency spectrum will be covered by the subjective description. The description will also include any other nuances that are too difficult to present using conventional graphs or other objective methods.

In this work, recordings of real instruments are compared with the synthesized sounds. These recordings are contained in a digital format from two online sources. The London Philharmonic Orchestra supplies a database with recordings of numerous orchestral instruments (*Sound Samples*, n.d.). Some instruments were not available in this database, which is why a second online database supplied by the Electronic Music Studios of the University of Iowa was used (Hallaron et al., 2012).

Two sources containing generic descriptions of synthesizer patches for specific instruments were used. Welsh's Synthesizer Cookbook (Welsh, 2006) presents very simple patches for a large variety of sounds of acoustic instruments. The author, Fred Welsh, assumes a preset interconnection of modules and only changes the settings of these modules to create the desired sound. The Sound on Sound magazine (*Sound on Sound*, n.d.) contains a series of articles titled Synth Secrets, in which the author Gordon Reid explains the design process of patches for a selection of instruments. For each sound, at least one patch from each source was set up and then recorded. In some cases, information from either source was used to produce a custom patch.

To record the sounds and their spectra, a suitable experimental setup was required. The output of the patch was sent into the line in port of a laptop. The JACK Audio Connection Kit (*JACK Audio Connection Kit*, n.d.) was then used to route this audio to a soundcard oscilloscope (Zeitnitz, 2017). Using this tool to compare the spectra of real instruments with the output of the synthesizer, the patch was built up and adjusted. Then, Audacity (Audacity Team, 2017) was used to record the patch outputting a tone at a frequency of 440 Hz, then one octave below that, and then outputting a chromatic scale from 220 Hz to 880 Hz. To create the spectrograms, the

Soundcard Oscilloscope was used to create averaged FFTs of each complete sound. The raw plot data was exported and then plotted using Gnuplot (*Gnuplot*, 2018), to be incorporated in this paper which was written using LaTeX (*The LaTeX Project*, n.d.). For each patch, the 440 Hz sound has been plotted for use in this work because it is a frequency that is well within the playable range of all instruments presented here. The bass drum however is an unpitched instrument, so no specific frequency could be specified for use in the FFT plots.

3 Instrument Synthesis

3.1 Orchestral flute

3.1.1 Analysis

The orchestral flute is in the woodwinds family, even though most flutes are made of metal. Like with members of the brass family, the flute's sound is produced by the oscillation of an air column. The column of air in the central cavity of the flute is excited by blowing into a hole positioned near one end of the instrument. Different pitches are produced by opening and closing holes along the length of the flute, which effectively changes the length of the air column. As with any enclosed air column, the wavelengths produced by the flute are theoretically multiples of the length of the air column. However, in the case of the flute, the wavelengths of higher frequencies are slightly longer than the flute. According to Reid (2003), this is because the air column is not excited at one end of the flute, but rather some distance from the end of the flute (Reid, 2003a). The overtones, which are approximately integer multiples of the fundamental frequency, make a flute sound very similar to a pulse wave. Yet, according to Reid (2003), due to limitations in the construction of a flute, there is a strong fall-off in harmonic amplitude above a frequency of about 2 kHz (Reid, 2003a). Thus, a raw pulse wave needs to be lowpass filtered to achieve the correct contour of the frequency spectrum. The sound of the flute is almost immediately at a maximum amplitude and stops very quickly as well. The tail of the note is slightly longer in duration than the beginning because it takes some time for all vibrations to dissipate.

3.1.2 Patches

Welsh (2006) uses a pulse wave with a pulse width of 25% sent through a lowpass filter. The cutoff frequency of the filter is modulated by an attack–decay envelope, causing the cutoff frequency to be increased at the beginning of the note and then to be decreased slightly (Welsh, 2006, p. 67). This simulates the way a flute player might blow with a varying pressure during the duration of the note as well as capturing the beginning and tail of the note, where the blowing pressure is built up or slowly removed by the player. The narrow pulse width of the oscillator removes single harmonics in the frequency spectrum, giving the sound a sharper, harsher quality. The amplitude of the sound is modulated slightly by an LFO, producing a tremolo effect. According to Reid (2003) tremolo on a flute would be achieved by varying the cross section of air passing into the flute. In contrast, vibrato would be achieved by varying the amount of air flow (Reid, 2003a). The envelope generator controlling the VCA is set to the same values as the envelope generator controlling the VCF. In effect this is a sharp 1/10s attack and a decay with approximately double that duration.

Reid (2003) proposes a similar setup, employing a filtered pulse wave oscillator. However, a pulse width of 50% is used, and a highpass filter is added after the lowpass filter with the aim of giving the final frequency spectrum a contour more like that of an actual flute. The highpass filter cutoff frequency is not modulated in any way. The cutoff frequency of the lowpass filter is modulated by an LFO to produce a vibrato–like effect. The amplitude of the vibrato is controlled by velocity CV, which means that a louder note will have stronger vibrato. In addition, lowpass filter cutoff frequency is offset by velocity CV, which has the effect of making notes brighter with increasing velocity. According to Reid (2003), this effect would be produced by increasing the blowing pressure for a real flute (Reid, 2003a).

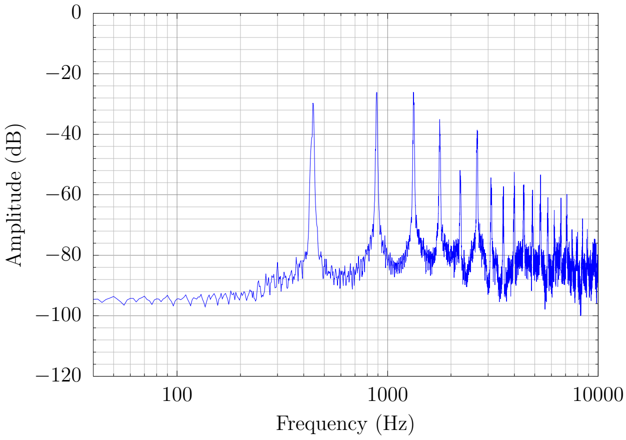

(a) The spectrum of a real flute, recorded by the Electronic Music Studios (Hallaron et al., 2012). Note the diminished amplitude of the first two peaks, and the dips in the spectrum.

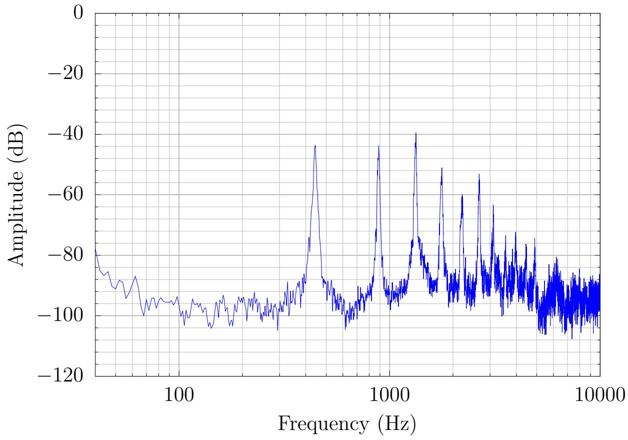

(b) The spectrum of a real flute, recorded by the London Philharmonic Orchestra (*Sound Samples*, n.d.). Note the diminshed amplitude of the first two peaks, and the dips in the spectrum.

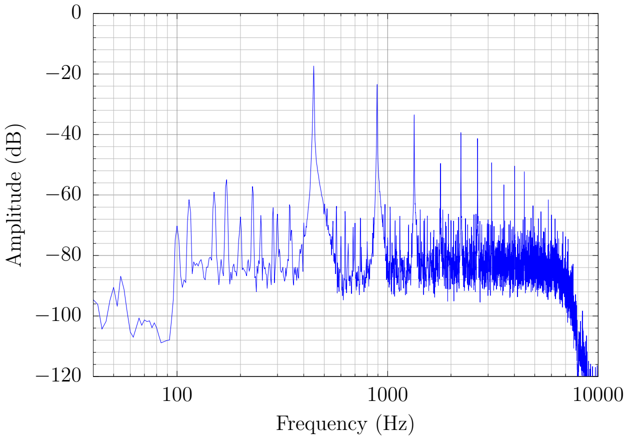

(c) The spectrum of Welsh (2006)'s flute patch (Welsh, 2006, p. 67). Dips in the frequency spectrum occur at 1.8 kHz, 3.3 kHz, and 5 kHz.

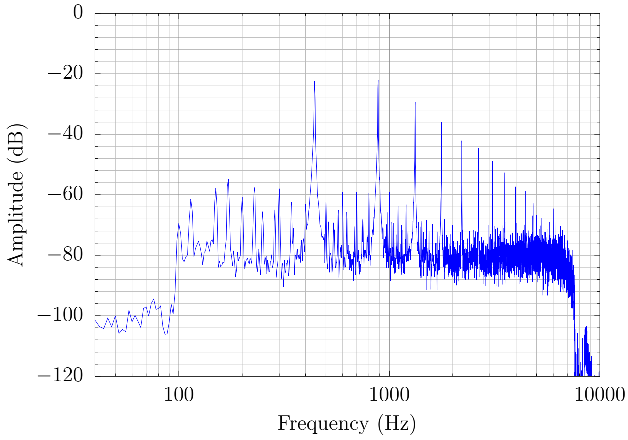

(d) The spectrum of Reid (2003)'s flute patch. The first peak has a diminished amplitude due to the use of a highpass filter.

Figure 4: FFT plots for the flute patch.

3.1.3 Conclusion

The patch presented by Welsh (2006) captures most of the timbre of an actual flute, but the articulation is very simple, which detracts from the realism of the sound. Also, the notes produced by the synthesizer sound too clean in the upper frequency ranges. According to Reid (2003), higher overtones in a flute are slightly off–pitch because the nodes of higher frequency waves extend over the end of the flute (Reid, 2003a). This is not simulated by the patch, as there is no simple way to modify single overtones with a synthesizer. Welsh(2006) has chosen to add tremolo to the flute sound (Welsh, 2006, p. 67) creating some interesting articulation, but this does not affect the frequency spectrum in any way other than modulating the amplitudes of the overtones slightly.

The patch presented by Reid (2003) extends the patch by Welsh (2006) by a single highpass filter which increases the realism of the timbre only slightly, because only low frequencies — which usually go unnoticed — are affected. The most important feature of Reid (2003)'s patch is, without doubt, the complex articulation, which allows the vibrato to swell and then slowly fade, similar to how a flutist might apply vibrato. The patch also gives the synthesizer player some control over the timbre of the sound by making the lowpass VCF cutoff frequency dependent on velocity. This emulates the different blowing strengths of a flutist to some extent. As stated previously, the blowing pressure affects mainly the timbre, not the volume of a flute. Similarly to the patch by Welsh (2006), this patch does not simulate the effect of off–pitch high frequency overtones. Reid (2003)'s patch nicely demonstrates how one setup can be used to emulate several similar instruments. Low pitched notes played on this patch sound very much like notes produced by members of the brass family, instruments such as the trumpet, tuba, or french horn. This is to be expected because such instruments produce sound by means of an oscillating air column in a similar fashion to that of the flute.

The frequency spectrum of the real flute is essentially a harmonic series were the amplitude of a peak is inversely proportional to its frequency. However, the first two peaks have a slightly lower amplitude than the third peak (See figure 4a and 4b). Reid (2003) attempted to simulate this effect via a highpass filter (Reid, 2003a), which decreases the amplitude of the first peak (See figure 4d). Welsh (2006) did not attempt to recreate this effect in any way (Welsh, 2006, p. 67). Some harmonics are somewhat less pronounced, creating dips in the spectrum, which Welsh (2006) simulated through the use of a 25% pulse width for the main VCO (Welsh, 2006, p. 67). Yet, the dips achieved by this method are too narrow and do not appear at exactly the correct frequencies in the spectrum (See figure 4c). Reid (2003) does not simulate this effect in any way. All in all, both patches produce a spectrum that is quite faithful to that of the flute, thus producing a moderately realistic sound.

3.2 Violin

3.2.1 Analysis

The Violin is a bowed chordophone. The sound of the violin is produced by rubbing a bow perpendicularly along one or more of the strings. The friction of the bow displaces and stretches the string until enough potential elastic energy builds up to overcome friction with the bow, and the string snaps back to its original position. This occurs repeatedly as the bow is dragged across the string and effectively causes the string to sustain a vibration. According to Reid (2003), the waves caused by this travel along the string and are transferred to the violin body via the bridge. The body of the violin removes and enhances some harmonics in a very complex manner, and a post connecting the top and bottom plates increases the rate of energy transmission (Reid, 2003b). The sound of the violin starts with less than 1/10th of a second, and decays for about a quarter to half second after the violin strings are no longer bowed.

3.2.2 Patches

Welsh (2006) uses a sawtooth oscillator synchronized to a master oscillator with a pitch six semitones lower than the slave. The slave oscillator waveform is then passed through a filter with a cutoff frequency at 1.9 kHz. The sound has an immediate attack and decays within less than a second. An LFO controlling the slave oscillator frequency adds some vibrato to the sound.

Reid (2003) proposes simply sending a sawtooth wave through a fixed filter bank to emulate the resonances of the violin body (Reid, 2003b). The envelope of the sound has a 1/10th second attack and a quarter second decay. Sometime after the note starts, vibrato is applied to the sawtooth oscillator. The patch was implemented using a filter bank which was set so that the synthesized frequency spectrum matched the frequency spectrum of a real violin as closely as possible. In an attempt to improve the sound of the patch, two settings for the fixed filter bank were tested.

3.2.3 Conclusion

Surprisingly, Welsh (2006)'s simple setup sounds very similar to Reid (2003)'s complex patch employing the fixed filter bank. However, the sound is very tinny and misses some of the low pitched characteristic sounds of a real violin. Also, the peaks of the frequency spectrum stay static throughout the duration of one note, which is very much unlike a real violin, where the spectrum changes over time as different vibrating parts of the violin exchange acoustic energy. Reid (2003)'s patch was tested using two different formant filter configurations. The first configuration shows more resemblance to a violin in the low frequencies than Welsh (2006)'s patch, and the complex vibrato creates the illusion that the frequency spectrum changes at the beginning of the sound. Note that the formant filter settings do not change in any way.

In the second formant filter configuration, the base notes were emphasized even more, so that this patch sounds somewhat like an oboe or clarinet. However, the high frequencies are by no means left out and contribute to the sound to make it sound very much like a violin as well. Again, the fact that the frequency spectrum does not change during the note detracts from the realism of the sound. Personally, I would not describe one of the configurations of the fixed filter bank as more realistic than the other. Instead, the two variations simply sound like two different violins. This is interesting in its own right because the filter settings were tweaked by comparing the vio-

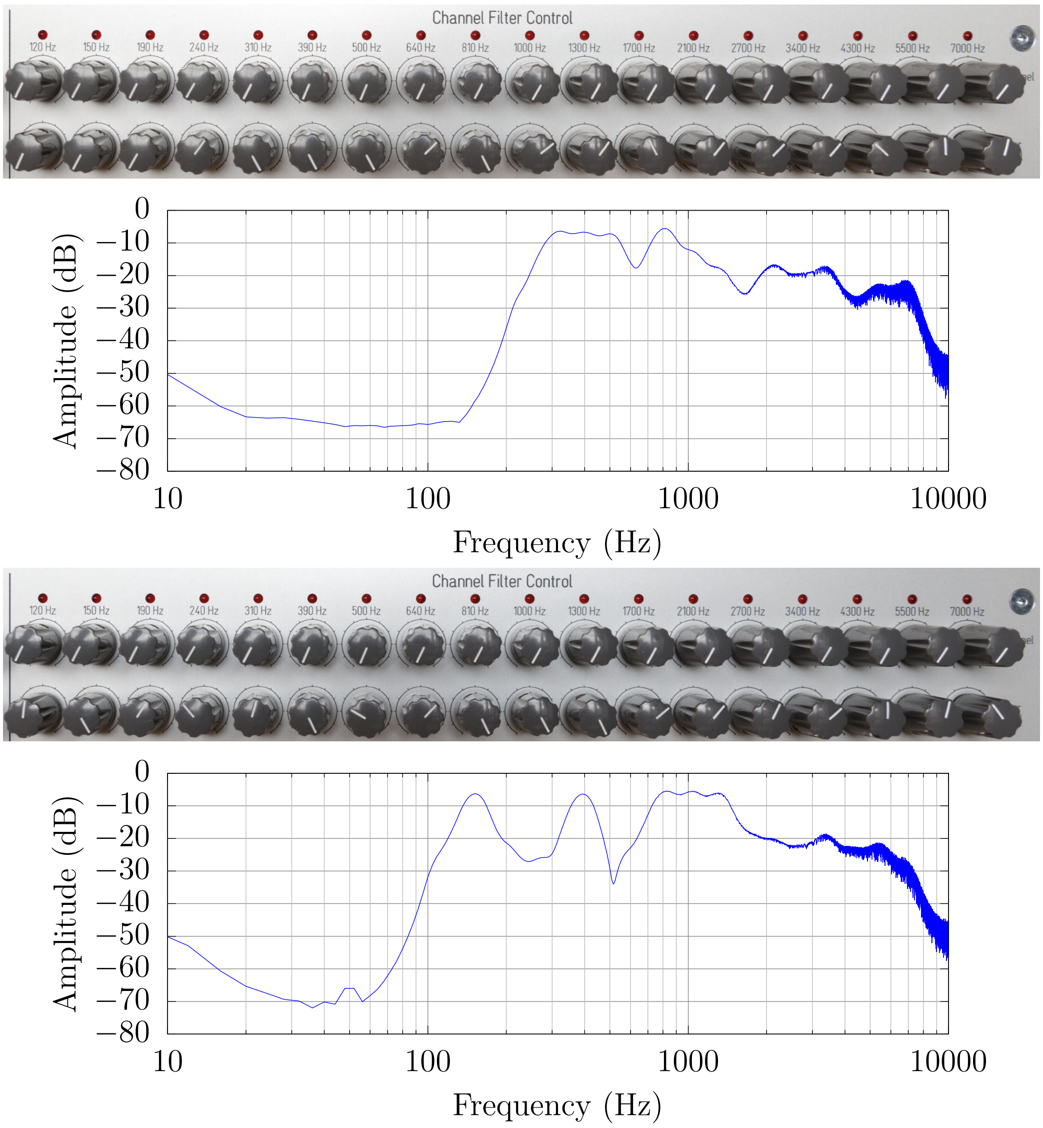

Figure 5: Two images showing the two settings used for the fixed filter bank. The bottom row of dials show the amplitude of sound at the frequency labelled above. The top row of dials has no effect in this setup. The first setting and FFT plot correspond to the first patch described.

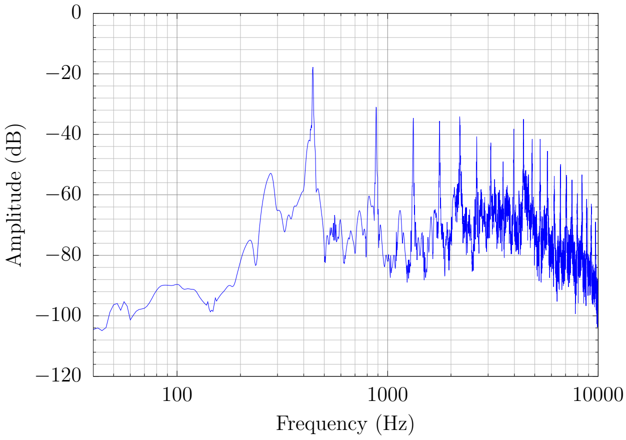

(a) The spectrum of a real violin, recorded by the Electronic Music Studios (Hallaron et al., 2012). There are dips at 1.4 kHz and 3.5 kHz. The amplitudes of the peaks higher than 4 kHz are proportional to the inverse frequency.

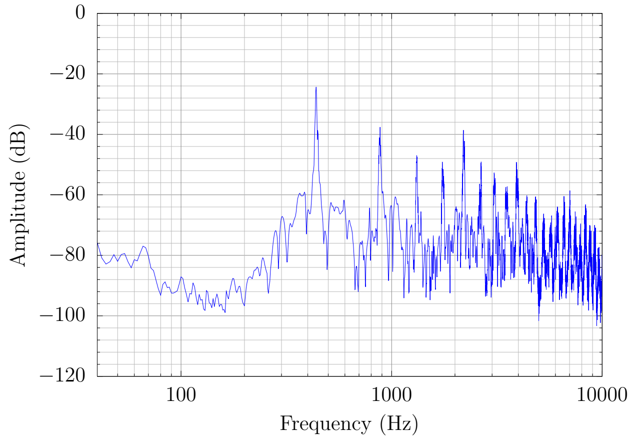

(b) The spectrum of a real violin, recorded by the London Philharmonic Orchestra (*Sound Samples*, n.d.). Dips in the spectrum are located at 1.9 kHz, 3.5 kHz, and 5 kHz.

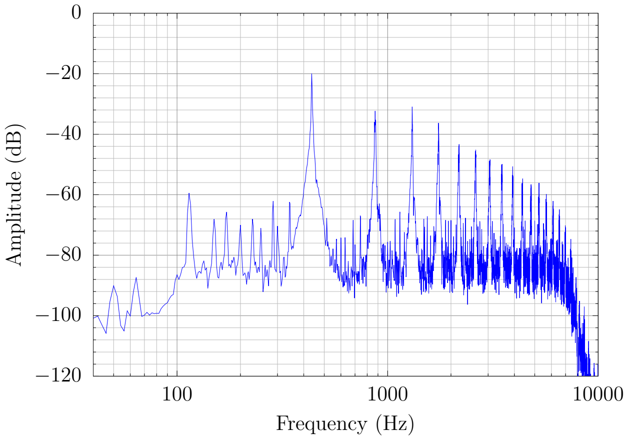

(c) The spectrum of Welsh (2006)'s violin patch. The spectrum is mostly regular, although the second peak has a diminished amplitude.

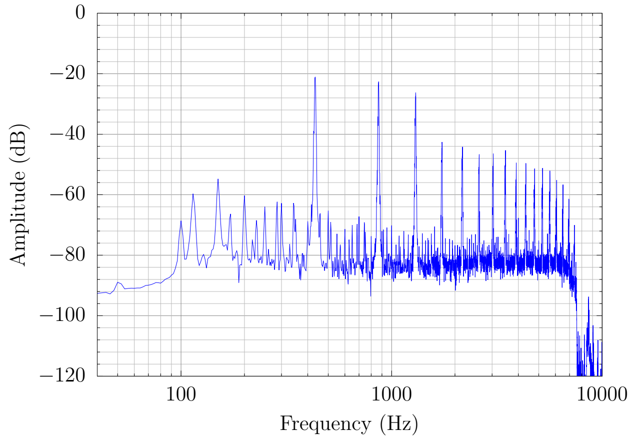

(d) The spectrum of Reid (2003)'s violin patch. The first three peaks have a larger amplitude than the rest.

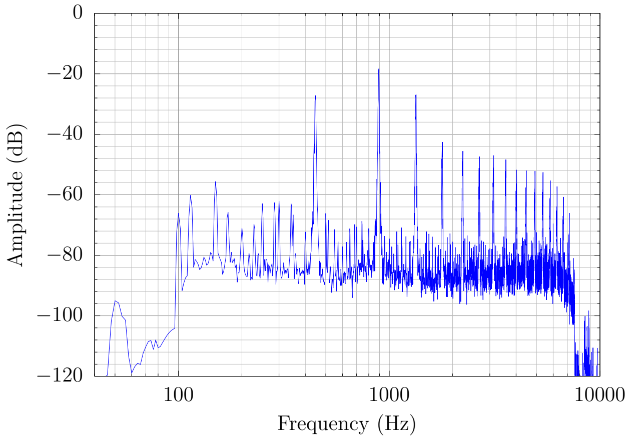

(e) The spectrum of Reid (2003)'s violin patch, with alternate settings for the fixed filter bank. Again, the first three peaks have a larger amplitude than the rest.

Figure 6: FFT plots for the violin patch.

lin patch to the same violin recording, which resulted in two violin–like but nonetheless different sounds.

In this case, the spectra of the two sample sources have some important differences. The spectrum of the Electronic Music Studios violin has a harmonic series of overtones, with two wide dips at 1.4 kHz and 3.5 kHz. After this, the amplitudes of the harmonics are inversely proportional to the frequency (See figure 6a). The spectrum of the London Philharmonic Orchestra has similar dips, but at different locations. One is located at 1.9 kHz, a second one at 3.5 kHz, and a third one at 5 kHz (See figure 6b). The spectrum of Welsh's patch is a normal harmonic series of frequency peaks, which means that the amplitude of a peak is inversely proportional to its frequency. The second harmonic has a slightly diminished amplitude (See figure 6c). For the first configuration of Reid's patch, we see another harmonic series. In this case, the first three peaks are larger than the other peaks, which only decrease very slowly in amplitude (See figure 6d). The second filter bank configuration has a frequency spectrum similar to the first configuration. Again, the first three peaks are raised and much more prominent than the rest. However, the second peak has the highest amplitude, being almost 7 dB in amplitude above the first and third peaks (See figure 6e).

3.3 Xylophone/Marimba

3.3.1 Analysis

The sound of a xylophone is produced by hitting a single wooden block suspended over a metal tube of a specific length with a mallet. According to Randel (1999), the block is shaped in a way to produce a specific pitch, which is then amplified by the cylindrical resonator tube beneath. The thickness and length of each block contribute to the resulting sound (Randel, 1999, p. 747). It is very difficult to physically analyze the vibrations produced by the block because it is a three–dimensional object that cannot be approximated to a simple geometrical shape like a string or membrane. The effect of the resonator tube is to enhance harmonics whose wavelengths correspond to multiples of the length of the tube. This turns the sound from the block without almost no discernible pitch into a sound with a well defined pitch. Due to these properties, the xylophone sound can be split into two separate parts, which are the noise of the vibrating wooden block and the ringing of the air column in the resonator tube.

3.3.2 Patches

Welsh (2006) proposes the usage of two separate triangle oscillators to mimic the ringing of the wooden xylophone block and the tap of the mallet hitting the bar. The oscillator creating the tap sound is set to a high frequency, while the other oscillator is tuned to produce the desired overall frequency of the sound. The mixed sound then passes through a lowpass filter without resonance with its cutoff frequency set to approximately 100 Hz. The cutoff frequency of the lowpass filter is increased by an envelope at the start of the sound to allow the passage of the mallet tap sound. The cutoff frequency then quickly returns to 100 Hz to allow only low frequencies to pass. An LFO adds some tremolo to the sound (Welsh, 2006, p. 84). Why Welsh has chosen to add tremolo is not clear, because the sound of the xylophone does not exhibit any such articulation. In the patch created for this work, sine wave oscillators were used instead of triangle oscillators, as the latter were not available. According to Vail (2014), sine wave oscillators have a similar sound to triangle oscillators, but they lack the distinct harmonic series (Vail, 2014, p. 11), which may have diminished the realism of the sound substantially.

Reid (2003) does not describe a specific patch for the xylophone but provides some general insights on percussion synthesis instead. Using these, a more complex patch was created as follows. A single sawtooth oscillator was sent through a lowpass filter to mimic the sound of the oscillating air column. The filter had an envelope applied to it, decreasing the cutoff frequency over time. Two detuned sawtooth oscillators were patched through a ring modulator to generate the plethora of inharmonic frequencies present in the tapping sound of the mallet. It would have been inappropriate to simply use an unfiltered white noise source, as a wooden block does not vibrate evenly at all frequencies. It may have been possible to use a noise source patched through a formant filter or fixed filter bank, but a ring oscillator produces approximately the correct frequencies and is much simpler to set up. The tapping sound was passed through a lowpass filter whose cutoff frequency was quickly lowered after the start of the sound to remove the mallet tap. This was achieved using an attack–decay envelope. The tapping and ringing sound were mixed together for the final audio signal.

3.3.3 Conclusion

The sound produced by Welsh (2006)'s patch is realistic to some extent but sounds too clear for a xylophone. It does not sound as if a wooden block is being struck, but rather some other material like glass or metal. The tremolo sounds unnatural but I decided to keep it because it distracts from the fact that the frequency spectrum is not very close to that of a xylophone. In Welsh's patch, the envelope creates articulation that sounds very accurate. The start of the note is soft and muted, after which the sound trails off in about half a second.

The custom sound is even more unnaturally clear and clean, and it is quite easy to hear that a sawtooth oscillator was used, even though it is passed through a filter. The sound also disappears too quickly, but setting a longer envelope decay time would be inappropriate too because it would make the characteristic sound of the sawtooth oscillator even more audible. Perhaps the decay curve of a standard ADSR envelope is simply the wrong shape for the purpose of simulating a xylophone, where the decay is very fast at first, but the sound sustains for a long time, albeit very quietly. In Welsh (2006)'s patch, this is simulated to some extent because some sound is leaked through the VCA due to the LFO tremolo modulation. This is because the LFO modulation signal is added to the VCA envelope signal which means that the sound is never completely muted.

The spectrum of a real xylophone consists of a single peak at 440 Hz with two sidebands at approximately 15 dB below the peak, located at 400 Hz and 500 Hz. Another, smaller peak appears at 1.3 kHz, and a larger one at 2.9 kHz. From this point on, there seems to be some kind of harmonic

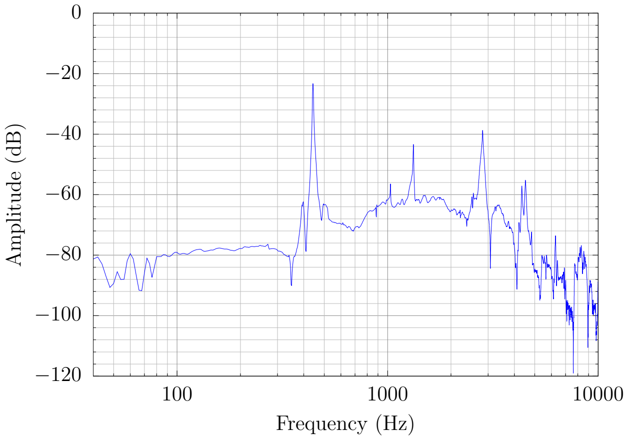

(a) The spectrum of a real xylophone, recorded by the Electronic Music Studios (Hallaron et al., 2012). Note the characteristic sidebands of the peak at 440 Hz.

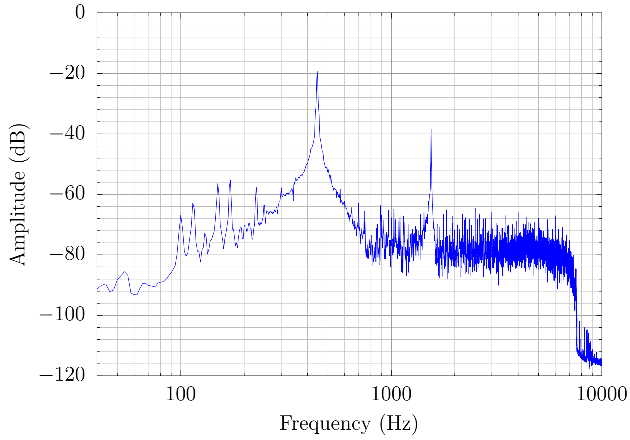

(b) The spectrum of Welsh (2006)'s xylophone patch. The peak at 440 Hz lacks the characteristic sidebands observed in the spectrum of a real xylophone, and is very wide and triangular instead.

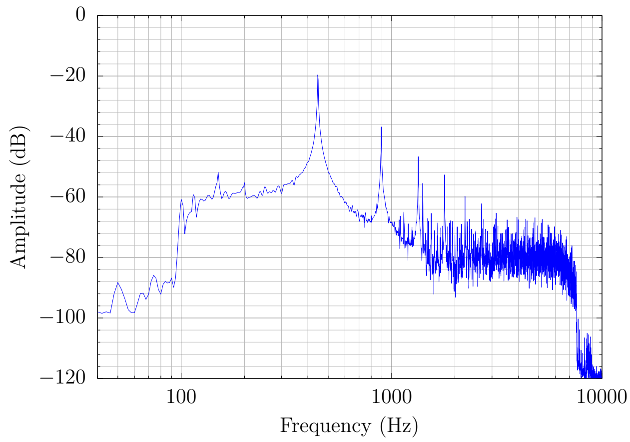

(c) The spectrum of the custom xylophone patch. Again, the first peak is wide and triangular. The harmonic series of the sawtooth oscillator is very prominent.

Figure 7: FFT plots for the xylophone patch.

series, with more peaks at 3.4 kHz and 5.2 kHz (See figure 7a). The patch by Welsh (2006) has a large, wide peak at 440 Hz. This peak does not have the sidebands observed in the real spectrum and instead has a very wide, triangular base. In the spectrum of Welsh (2006)'s patch, a second peak appears at a frequency slightly higher than the frequency of the second peak of the real spectrum, at approximately 1.6 kHz instead of 1.4 kHz. After this, no more peaks can be observed (See figure 7b). The custom patch's first peak looks just like the first peak of Welsh's patch. However, after this, the harmonic series of the sawtooth oscillator becomes very prominent, and the spectrum is very different from that of a real xylophone, which has only very few peaks.

3.4 Guitar

3.4.1 Analysis

The Guitar produces sound via the vibration of strings of varying thickness stretched across a hollow body that amplifies the string vibrations and shapes the frequency spectrum of each note. According to Reid (2001), the theoretical spectrum produced by a vibrating string is that of a triangle wave (Reid, 2001b). However, the final sound produced by the guitar is by no means similar to a simple triangle wave. The vibrations of a single string are transferred to the body and the air inside, where some harmonics are enhanced and others suppressed. The other five strings also start vibrating and contribute to a complex mix of many vibrating physical bodies that interact with one another and exchange energy in different parts of the frequency spectrum (Reid, 2001b). Without doubt, the sound of a guitar can at most be approximated by an analog or even digital synthesizer.

3.4.2 Patches

Welsh (2006) uses two pulse wave oscillators with the second oscillator detuned to 10 semitones above, and an amplitude of 4 dB less than the first oscillator. The first oscillator uses a pulse width of 25% , while the second uses a pulse width of 10% (Welsh, 2006, p. 55). The different pulse widths of the oscillator waveforms introduce narrow dips into the spectrum much like those found in the frequency spectrum of a real guitar, where they are a result of the destructive interference caused by the interactions of several strings with each other and the body of the guitar. The sound is passed through a filter with a cutoff frequency of 2.0 kHz, which is increased at the beginning of the sound using an envelope to replicate the sound of plucking the string. The cutoff frequency then slowly decreases, causing the sound to slowly become lower and duller, and finally die away altogether.

Reid (2001) proposes a theoretical patch for the guitar, employing more than 100 modules and several filters that have an impossibly high attenuation slope (Reid, 2001c). The theoretical patch synthesizes each of the 6 strings of a guitar separately and attempts to give the synthesizer player separate control over every single note produced. A guitar player has the possibility to control the timbre of every single note separately as well. However, the resulting patch is not only unplayable but also unbuildable, at the very least due to the physically unachievable specifications of the filters. In a second article, Reid (2001) tries to simplify this ultimate patch by focusing on the synthesis an electric guitar but concludes that the original theoretical patch does not have any redundant or unneeded parts. Reid (2001) then continues by presenting and analyzing the solutions that some synthesizer companies have used for some of the patch presets of their synthesizers (Reid, 2001a). The patches described by Reid (2001) have been analyzed and the most important parts grafted to create a custom patch. Two different oscillators in unison — one using a pulse wave with a pulse width of approximately 20% and the other using a sawtooth wave — were fed through a filter with much the same parameters as the patch by Welsh (2006). One of the patches presented by Reid (2001) changed the pulse width over time to introduce some change of timbre as the sound trailed off, but this feature made the custom patch sound too much like an electric guitar and was thus left out. However, this is another example of how the patch could be modified to sound like a similar instrument.

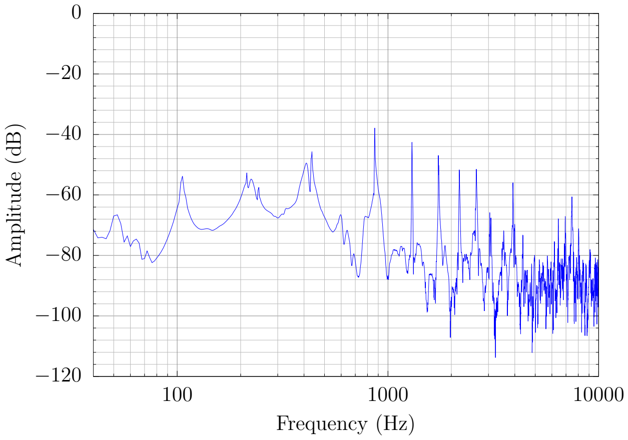

(a) The spectrum of a real guitar, recorded by the Electronic Music Studios (Hallaron et al., 2012). Note the low amplitude and strange shape of the first peak.

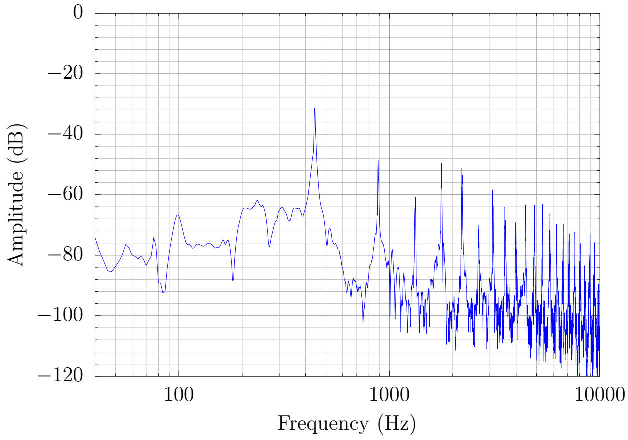

(b) The spectrum of a real guitar, recorded by the London Philharmonic Orchestra (*Sound Samples*, n.d.). Note the differences to the recording done by the Electronic Music Studios. The amplitudes of the overtones are more irregular, and there is more spectral content above 1.3 kHz than in figure 8a.

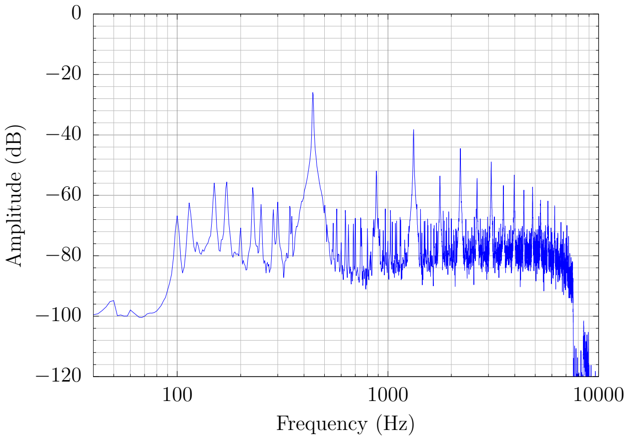

(c) The spectrum of Welsh (2006)'s guitar patch. The amplitude of every second harmonic is slightly diminished, making this a very regular spectrum. None of the real instrument spectra exhibit this much regularity.

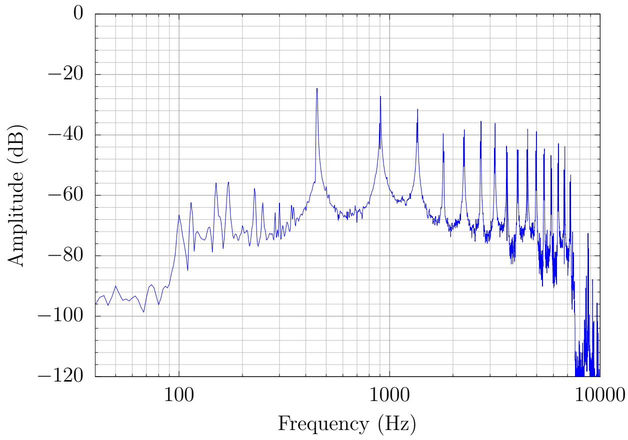

(d) The spectrum of the grafted guitar patch. The amplitude dips in this spectrum are much more irregular than those in figure 8c.

Figure 8: FFT plots for the guitar patch.

3.4.3 Conclusion

The patch described by Welsh (2006) produces a very dry sound that seems to lack some overtones compared to a guitar. The sound dies off rather quickly — not in amplitude, but in a harmonic sense — losing most of its upper frequencies soon due to the sweep of the lowpass filter. Otherwise, the timbre of the note produced is rather harsh and would be attributed to the higher–pitched strings of the guitar, which have a harsher timbre than the lower–pitched strings, whose sound is slightly more sonorous.

The custom patch produces a far more direct timbre, in the sense that lighter filtering allows more overtones to appear in the final sound. However, this also makes it much more audible that a very basic waveform was used, which detracts from the sound's credibility. The filter cutoff modulation in this sound was far more realistic than that in Reid's patch. In this patch, the sound wavered, highlighting different harmonics before dying away which is much more reminiscent of a guitar than Reid's patch, where high harmonics simply fade away in a monotonous manner. Both patches sound similar to a steel string or electric guitar, in contrast to a concert guitar, whose sound is slightly sweeter.

One observation made from the plots of the two different samples of a real guitar is that the spectra look very different. This is surprising because the spectra of the other instruments were in general more similar to each other. In the Electronic Music Studios spectrum, there is a series of increasingly quieter harmonics beginning at approximately 870 Hz. The fundamental frequency at 440 Hz forms quite a wide peak (See figure 8a). In the London Philharmonic Orchestra spectrum, the fundamental frequency peak is wide as well, but also louder. There are more dips in the amplitudes of the harmonics, and there is a lot more harmonic content above 1.3 kHz than in the Electronic Music Studios spectrum (See figure 8b). The Electronic Music Studios recording sounds like a steel string guitar, while the London Philharmonic Orchestra recording sounds a lot more like a concert guitar. The patch by Welsh shows some similarities to both spectra, having a wide fundamental frequency peak. Every second harmonic is slightly softer than its neighbors (See figure 8c).

The holes present in the London Philharmonic Orchestra spectrum are not as frequent or regular as this. For the custom patch, the fundamental peak is by far not as wide as in the other frequency spectra. The spectrum is quite similar to that of the London Philharmonic Orchestra recording, containing some irregular dips, and many tightly packed frequency peaks above 1.3 kHz (See figure 8d).

3.5 Bass drum

3.5.1 Analysis

The bass drum is an untuned membranophone. It consists of two parallel, circular membranes stretched over both ends of a short, wide cylinder. One of the diaphragms is hit using a foot–operated mallet. When this happens, it starts to vibrate in a complex way, typical of a thin physical sheet. These vibrations are then passed through the air inside the drum and the cylinder walls to the second membrane. This also starts vibrating, and the vibrations from both membranes travel back and forth and interact in complex ways. Sometimes, the second, unagitated diaphragm has a hole, presumably to allow better passage of vibrations to the outside of the drum.

3.5.2 Patches

For this patch, Welsh (2006) simply mixes two triangle waves in unison to thicken the peaks in the frequency spectrum and applies an envelope with an instantaneous attack and quarter second decay to the VCA. In the test patch, sine wave oscillators were used instead of triangle oscillators, as the latter were not available. According to Vail (2014), sine wave oscillators sound quite similar to triangle oscillators (Vail, 2014, p. 11).

Reid (2002) proposes a rather complex patch employing a sawtooth oscillator patched through a lowpass filter for the actual sound of the drum membrane. A sawtooth oscillator frequency modulated by a second sawtooth oscillator is passed through a succession of low- and highpass filters to create the tap of the mallet on the membrane. The cutoff frequency of the highpass filter is raised at the start of the sound to let the tap sound through and is then lowered below the cutoff frequency of the highpass filter to effectively mute the mallet sound. Both audio signals are then simply mixed together to form the final sound. (Reid, 2002b) Instead of using a succession of lowpass and highpass filters which were unavailable, a bandpass filter was used for the mallet sound in the test patch. When the cutoff frequencies of a lowpass and highpass filter in series are very close together, the effect of the two filters can be approximated by a single bandpass filter. After the start of the sound, the bandpass cutoff frequency was decreased far enough to effectively mute the mallet sound.

Surprisingly, this was the only patch where Reid and Welsh did not agree on any of the elements of their respective patches. This may have arisen from the fact that Reid relies on physical analysis, while Welsh attempts to find the best settings for a generic modular synthesizer system.

A third patch was tested due to the unsatisfactory nature of the previous patches. This third patch is a preset on the Arp Axxe synthesizer and is described in (Reid, 2002a). The patch uses three basic sound sources, consisting of a pulse wave oscillator, a sawtooth oscillator, and white noise. These sources are then mixed and passed through a lowpass filter with maximum resonance, causing self–oscillation, which means that the filter produces an approximate sine wave at the cutoff frequency. The filter is set to the lowest possible cutoff frequency. This is then controlled by an envelope to produce the bass drum sound. While the proposed patch uses two outputs of the same oscillator to produce pulse and sawtooth waves, two different oscillators were used in the test patch. However, the difference in the final sound is minimal because the oscillators were tuned to be exactly in unison at the needed pitch.

3.5.3 Conclusion

The patch proposed by Welsh sounds more like a low pitched bongo drum than a bass drum in the sense that a clear pitch is discernible, and the sound is sweeter and more harmonic than the untamed noise of a bass drum. It also

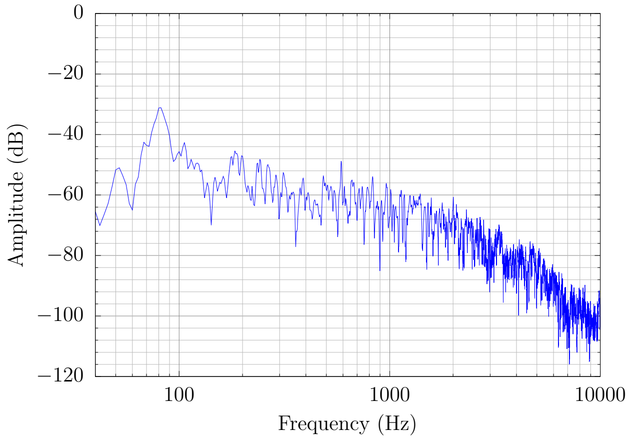

(a) The spectrum of a real bass drum, recorded by the Electronic Music Studios (Hallaron et al., 2012). There is one wide peak at 80 Hz. Noise with an amplitude that is louder than usual persists until approximately 4 kHz.

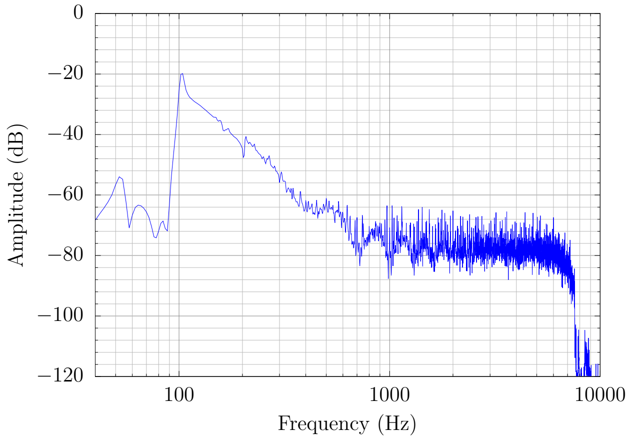

(b) The spectrum of Welsh (2006)'s bass drum patch. The only peak has a very strange triangular shape between 90 Hz and 700 Hz.

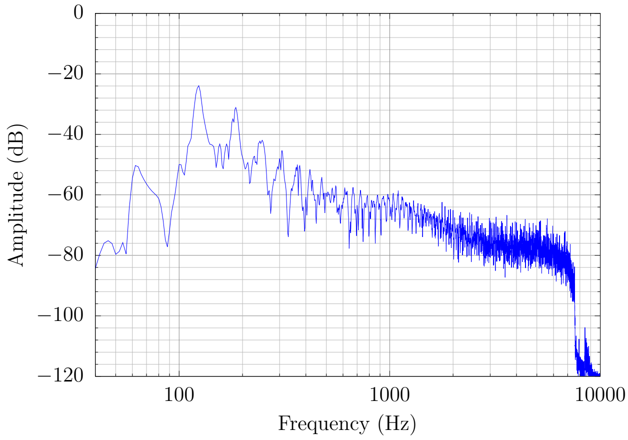

(c) The spectrum of Reid (2002)'s bass drum patch. The weak, but still very clear progression of overtones makes this spectrum very different from the spectrum of the real bass drum in figure 9a.

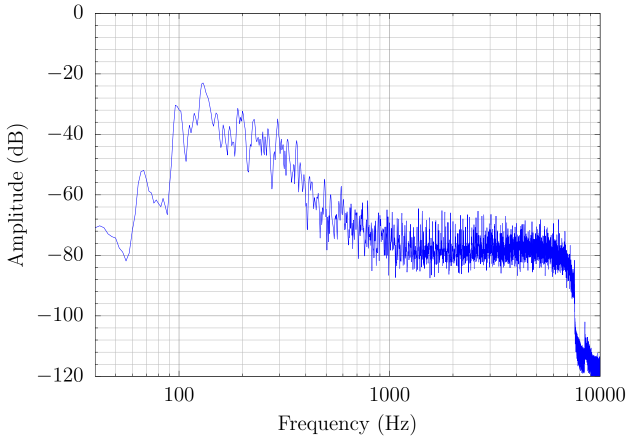

(d) The spectrum of the Arp Axxe synthesizer bass drum patch. The peak in this spectrum is very similar to the peak in the spectrum of the real bass drum in figure 9a. However, the amplitude of the noise at higher frequencies falls off too quickly.

Figure 9: FFT plots for the bass drum patch.

seems like the attack of the envelope is too short. However, using a longer attack would result in an audible amplitude ramp, which is also undesirable. Perhaps the filter should be controlled by an envelope as well. The attack of this envelope would be slightly longer than the attack of the amplitude envelope to simulate the time it takes for the vibrations of the first membrane to transfer to the second membrane.

The patch by Reid has an even better defined pitch that decreases over the duration of the note. This alteration of pitch sounds unrealistic because such a change in pitch does not occur for a bass drum, which is completely unpitched in reality. There is some low–pitched oscillation that makes the drum patch sound as if an object was held close enough to the drum membrane to touch it as it was vibrating. The result is a rattling sound, which may be more realistic for a different type of idiophone. Here, they are unfitting.

The third patch is by far the best and most realistic sounding. It manages to capture what sounds like the whoosh of air caused by very low–pitched vibrations of the drum membranes. Although the sound of the patch is not completely unpitched, this is concealed by the fact that the pitch is very low and does not change as the note develops. If the patch is played at higher frequencies, the pitch becomes noticeable, and the patch sounds similar to a timpani or kettle drum. As was the case with Reid's setup, there is some very low amplitude modulation that produces a rattling sound, but this is much less prominent than in Reid's patch and is more fitting.

The real bass drum spectrum has a very wide peak at approximately 80 Hz. At higher frequencies, there is simply noise, decreasing in amplitude as the frequency increases (See figure 9a). For Welsh's patch, there is a very prominent triangular peak at 100 Hz. The left flank is very steep, while the right flank decreases proportionally to the inverse frequency until about 700 Hz, from which point on the spectrum is simply very quiet noise with no change in amplitude (See figure 9b). The patch by Reid (2002) shows a clear progression of harmonic overtones beginning at approximately 120 Hz, which explains why Reid (2002)'s patch sounds so unrealistic. There are no notable harmonics present in the spectrum of a real bass drum (See figure 9c). The patch originally designed for the Arp Axxe synthesizer has a peak around

130 Hz which looks very similar to that of the real bass drum, featuring several smaller sidebands. After this, there is some noise which descends to a minimum proportionally to the inverse frequency until approximately 1 kHz (See figure 9d).

4 Conclusion

In this work, patches for five different generic instruments were programmed on an analog modular synthesizer. Most of the patches were taken from existing literature, while one of each of the xylophone, guitar, and bass drum patches were either designed by myself or were the result of combining several patches from various sources. Some patches had to be modified to be compatible with the available hardware, but presumptively, this did not have a noticeable impact on the sound. Each patch was recorded and the frequency spectrum, amplitude envelope, and personal impression were used to analyze the resulting sound. These metrics made it possible to compare and contrast the positive and negative features of the various patches for each sound.

It has been shown that some groups of instruments lend themselves better to electronic synthesis than others. For example, very satisfying patches can be produced for aerophones, while membranophones such as the bass drum are slightly more difficult. In the large category of idiophones, I chose to synthesize the xylophone. While none of the patches for the xylophone were completely satisfactory, I believe that other instruments of the idiophones group may be easier to synthesize. This could be the topic of further research. The fourth group of instruments, the chordophones, proved to be very difficult to synthesize. Despite the use of a very large fixed filter bank, the violin patches were not very realistic. The synthesized guitar was even less satisfactory. Perhaps the large number of interacting physical components makes the sounds of these instruments too complex for analog synthesis. A flute can be simplified to an oscillating air column, which makes it easier to synthesize than a guitar, where the vibrations of physical three–dimensional objects have to be taken into account.

An interesting observation is that in all cases, both Welsh's and Reid's patches competed closely in terms of sound realism. This is somewhat surprising because in contrast to Reid, Welsh optimizes patches for very simple synthesizers with a fixed architecture. I was not expecting this personally, and it shows that complexity is not necessary for creating very useful and realistic sounds.

List of Figures

References

Audacity Team. (2017). *Audacity(R).* `https://www.audacityteam.org/`. (Retrieved 00:47, February 7, 2018)

Ekbert Hering, M. S., Rolf Martin. (1989). *Physik für ingenieure* (3rd ed.). VDI Verlag.

Gnuplot. (2018). `http://www.gnuplot.info/`. (Retrieved 19:59, February 14, 2018)

Hallaron, M., Adamson, S., Durian, E., Ritz, J., Nichols, J., Cash, M., ... Frantz, D. (2012). *Musical instrument samples.* `http://theremin.music.uiowa.edu/MIS.html`. (Retrieved 22:05, December 12, 2017)

Jack audio connection kit. (n.d.). `http://jackaudio.org/`. (Retrieved 17:49, December 29, 2017)

Paul A. Tipler, G. M. (2004). *Physics for scientists and engineers* (5th ed.). W.H. Freeman and Company.

The LaTeX project. (n.d.). `https://www.latex-project.org/privacy/`. (Retrieved 20:05, February 14, 2018)

Randel, D. M. (1999). *The harvard concise dictionary of music and musicians.* The Belknap Press of Harvard University Press.

Reid, G. (2001a). *A final attempt to synthesize guitars.* https://web.archive.org/web/20160403120425/http://www.soundonsound.com/sos/Oct01/articles/synthsecrets30.asp. (Retrieved 12:35, December 3, 2017)

Reid, G. (2001b). *Synthesizing plucked strings.* https://web.archive.org/web/20160404062927/http://www.soundonsound.com/sos/Aug01/articles/synthsecrets28.asp. (Retrieved 12:33, December 3, 2017)

Reid, G. (2001c). *The theoretical acoustic guitar patch.* https://web.archive.org/web/20160404062915/http://www.soundonsound.com/sos/Sep01/articles/synthsecrets29.asp. (Retrieved 12:34, December 3, 2017)

Reid, G. (2002a). *Practical bass drum synthesis.* https://web.archive.org/web/20160403123708/http://www.soundonsound.com/sos/Feb02/articles/synthsecrets0202.asp. (Retrieved 12:37, December 3, 2017)

Reid, G. (2002b). *Synthesizing drums: The bass drum.* https://web.archive.org/web/20160405040308/http://www.soundonsound.com:80/sos/jan02/articles/synthsecrets0102.asp. (Retrieved 12:36, December 3, 2017)

Reid, G. (2003a). *Practical flute synthesis.* https://www.soundonsound.com/techniques/practical-flute-synthesis. (Retrieved 12:19, December 3, 2017)

Reid, G. (2003b). *Synthesizing bowed strings: the violin family.* https://web.archive.org/web/20160405034337/http://www.soundonsound.com:80/sos/apr03/articles/synthsecrets48.asp. (Retrieved 12:21, December 3, 2017)

Sound on sound. (n.d.). https://www.soundonsound.com/. (Retrieved 18:17, January 3, 2018)

Sound samples. (n.d.). http://www.philharmonia.co.uk/explore/sound_samples/. (Retrieved 22:07, December 12, 2017)

Vail, M. (2014). *The synthesizer.* Oxford University Press.

Welsh, F. (2006). *Welsh's synthesizer cookbook* (3rd ed.). Author.

Zeitnitz, C. (2017). *Soundcard oscilloscope.* https://www.zeitnitz.eu/scope_en. (Retrieved 17:47, December 29, 2017)

Ziegenrücker, W. (2009). *Abc musik.* Breitkopf & Härtel.