WELSH'S SYNTHESIZER COOKBOOK

By Fred Welsh

3d Edition

Welsh's Synthesizer Cookbook

Dedication

This book is dedicated to my lovely wife Danielle for all her support and for putting up with all my late nights on the keyboards.

Many thanks to MDSP at Smartelectronix for giving permission to use Fre(a)koskope to produce the harmonic diagrams in the book.

CONTENTS

PATCHES

OSCILLATORS

Harmonics and oscillator waveforms

The most fundamental sound is the sine wave. Every sound whether it is natural or synthetic is made up of sine waves.

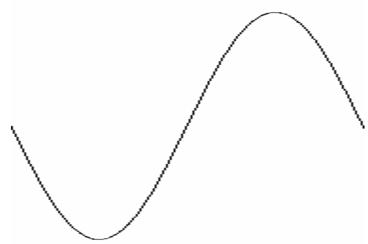

Figure 1. All sounds, natural and synthetic, are made up of sine waves

The sound of a sine wave could be described as extremely mellow and not particularly interesting, but when numerous sine waves of different pitches and amplitudes are mixed together the sonic possibilities become endless.

It is possible to create *any* waveform using nothing but sine waves. Let's begin by creating a sawtooth wave. Each sine wave is going to have a different pitch and as they go higher in pitch they will decrease in amplitude. The first one will have the lowest pitch and it will in fact determine the overall pitch of the sawtooth. This is called the fundamental frequency. For the case of this example we will give it a pitch of 440 Hz corresponding to the A4 key.

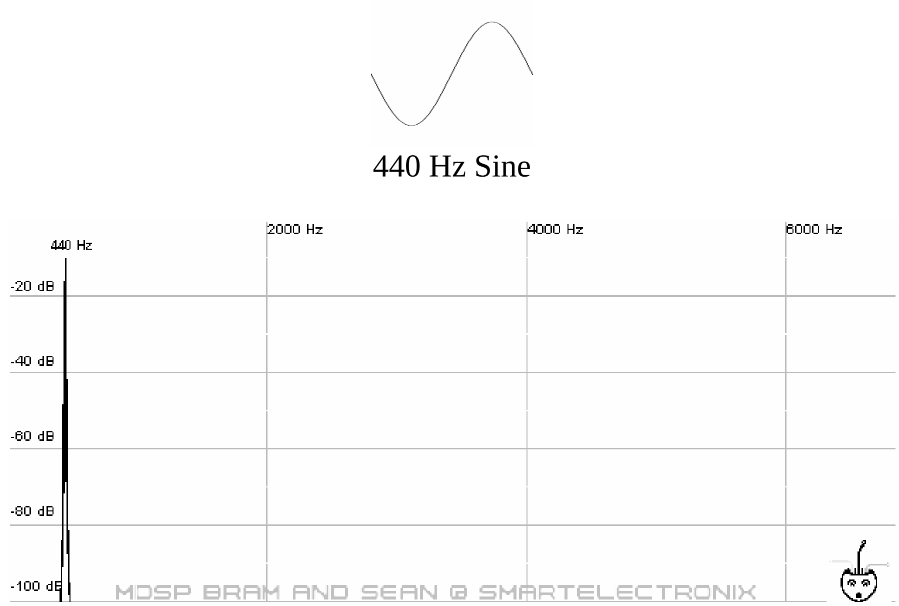

Figure 2. A 440 Hz sine wave and a graph representing its frequency and amplitude

The bottom portion of the figure is what is called a harmonic diagram. The horizontal axis is frequency/pitch and the vertical axis is amplitude. The spike represents the frequency and amplitude of the 440 sine wave. Let's add a second sine wave with a frequency of 880 Hz (440Hz x 2).

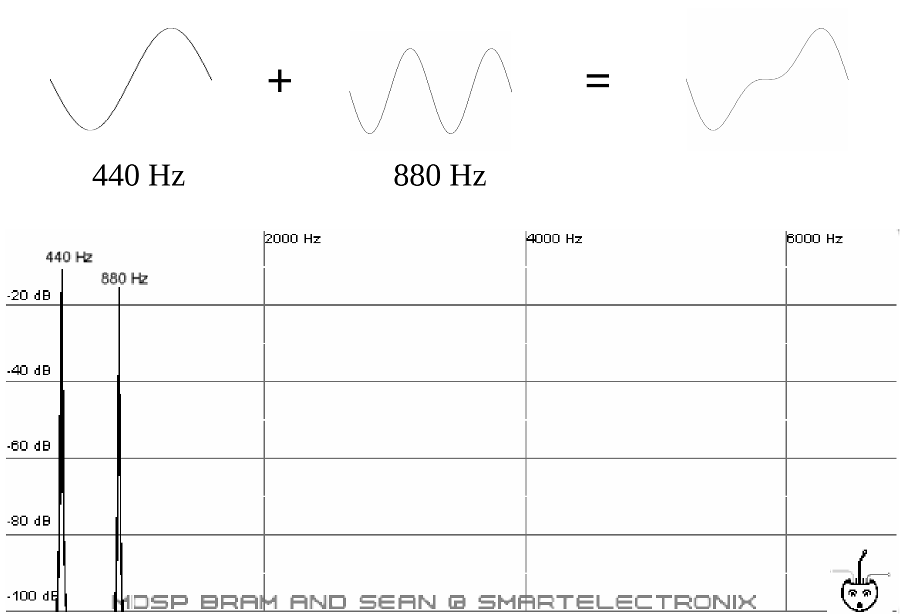

Figure 3. A 440 Hz sine wave added to a 880 Hz sine wave create a new waveform

Notice there is a second spike in the diagram representing the 880 Hz sine wave. Now add a third sine wave at a frequency of 1320 Hz (440 Hz x 3)

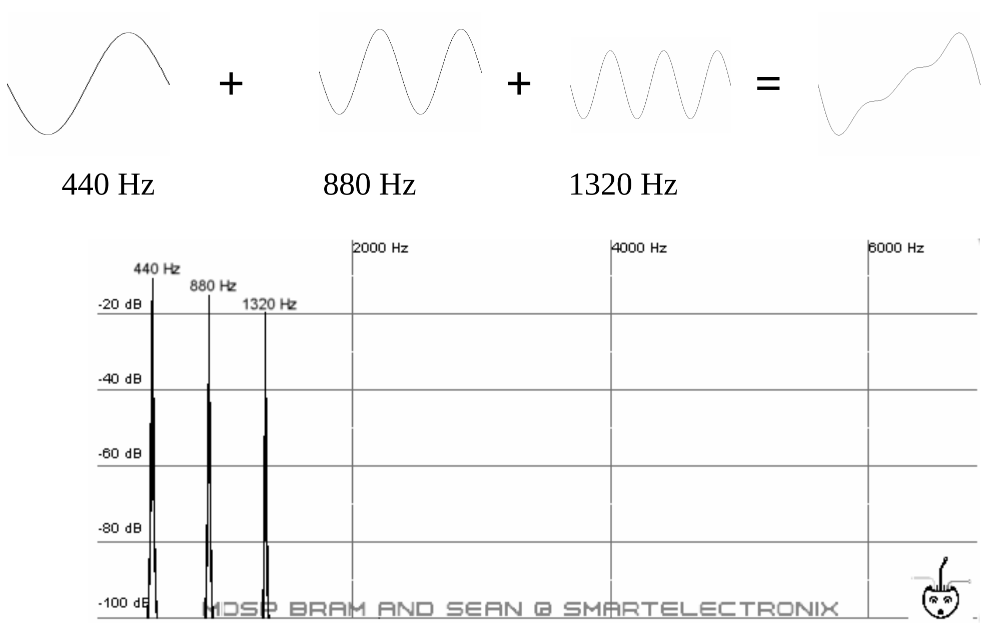

Figure 4. Three sine waves of frequency 440 Hz, 880 Hz, and 1320 Hz added together create a waveform that begins to look like a sawtooth wave.

Already the new waveform is beginning to lean to one like a sawtooth wave. Keep adding sine waves that are frequency multiples of 440 Hz and eventually a sawtooth wave is produced.

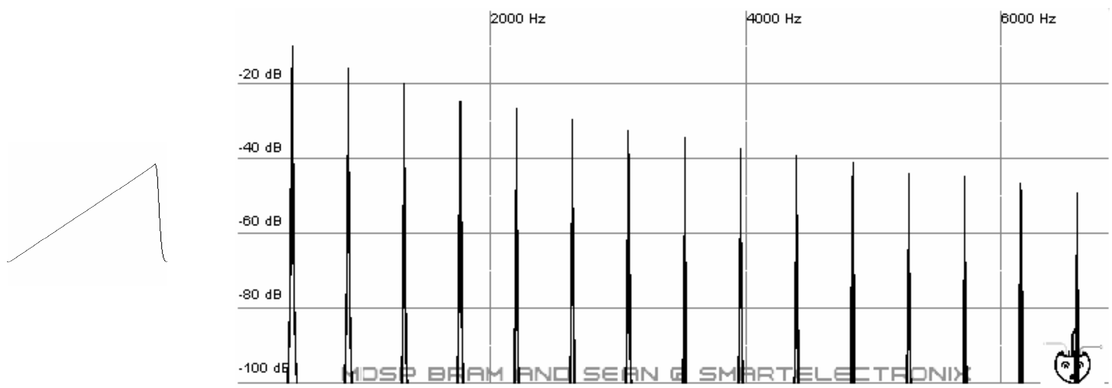

Figure 5. A sawtooth wave and its harmonic diagram

Each one of the harmonics(spikes) in the diagrams above are called partials. The first partial at the far left, as stated earlier, is known as the fundamental and it determines the pitch of the waveform. The other partials are all frequency intervals of the fundamental and they are known as overtones. Also notice that the fundamental is the loudest sine wave of all the harmonics and that as the overtones get higher in pitch their amplitudes decrease. There is of course nothing preventing us from mixing sine waves that all have the same amplitude it’s just that in this particular case it is a requirement of creating a sawtooth that the amplitudes diminish with higher frequency. Just as a point of interest if we were to mix these harmonics all at the same amplitude it would produce something that as successively higher-pitched sine waves are added would at first look like a sawtooth wave with a bad complexion and as more harmonics are added would look nothing like a sawtooth at all. Harmonic amplitudes are important!

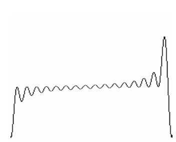

Figure 6. Waveform produced by setting the first 16 harmonics to equal amplitudes.

A square wave can be created using the same process. The difference is that a square wave is only made up of odd-numbered partials.

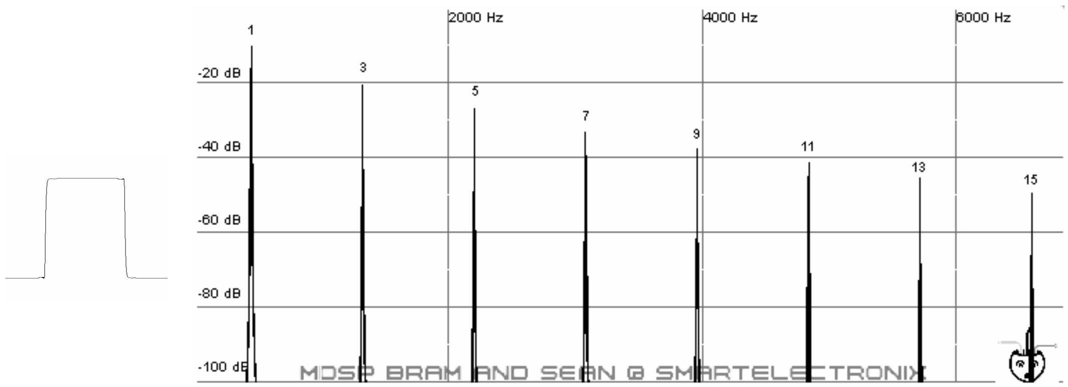

Figure 7. A square wave has only odd partials. This square wave was produced by an *actual* oscillator. Notice the slight imperfections in the waveform. No oscillator produces perfect waveforms nor do any two oscillators produce identical waveforms.

Triangle waves are also made up of odd-numbered partials but the higher harmonics are weaker than in a square wave.

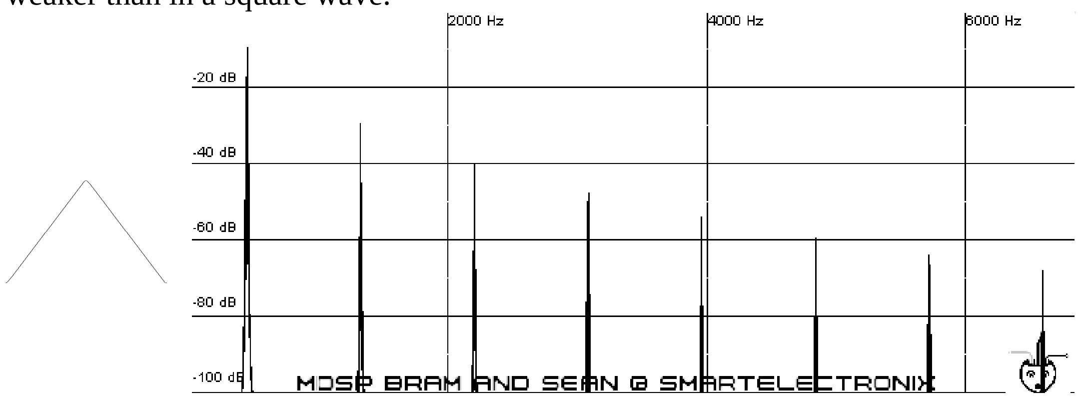

Figure 8. A triangle wave has odd harmonics that are weaker than those of a square wave.

The process of adding sine waves together to create sounds is known as *additive synthesis*. This method is used on a few digital synthesizers and soft synths and has also been used by pipe organs for hundreds of years. In a pipe organ each pipe produces a sine wave of a different pitch and by controlling the amount of air to each pipe it is possible to control the individual amplitudes of each sine wave which in turn makes it possible to produce sounds that are harmonically similar to other instruments. It could be said that pipe organs were the first synthesizers.

Analog synthesizers use a process called *subtractive synthesis* which is simply additive synthesis in reverse. Here's some synth terminonlogy for you: Sounds created by synthesizers are referred to as patches. This goes back to the early days of modular synthesizers when patch cables were used to route signals from module to module to create a sound. Patches created using subtractive synthesis start with waveforms that are already rich in harmonics such as sawtooth, square, and triangle waves. These waveforms are then passed to a filter which remove harmonics from the waveforms in order to produce the desired sound. The harmonics are *subtracted* out hence the process is known as subtractive synthesis.

The most common waveforms on analog synthesizers are the sawtooth, square/pulse, and triangle wave. There are three main reasons these are used almost universally rather than other waveforms. First is the fact that they all have lots of harmonics which can be chiseled away by the filter. Secondly they are relatively easy to produce using analog circuits. Thirdly they are each harmonically similar to broad families of acoustic sounds even without any filtering. This almost sounds as if to imply that acoustic imitation is the goal of synthesis which it most certainly is not. It is a practical criteria to start with if nothing else. Sawtooth waves are similar to brass and string instruments, square waves are similar to woodwinds, and triangle waves with their diminished higher harmonics are good for mixing together to produce inharmonic sounds (bells, chimes, etc) as well as adding the occasional rogue harmonic to saw and pulse waveforms. This may be surprising but even though these waveforms have similar harmonics and sound similar to certain families of instruments the waveforms of acoustic instruments look nothing like squares/pulses, sawtooths, or triangles. A waveform's usefulness lies with the harmonics that make it up and is NOT inherent to the shape of the waveform. Two sounds can have very similar harmonics yet have waveforms that look nothing alike. Adding waveforms together in an attempt to match the waveshape of an acoustic instrument won't produce desirable results most of the time.

While some synthesizers only have one oscillator and a few others have many, the majority have two oscillators that can be independently set to different waveforms. Two-oscillator configurations are so common because they are cheaper to produce than multi-oscillator synths yet they are much more capable of producing rich sounds than single-oscillator instruments. They are a good balance of economy and ability. By using two oscillators it is possible to not only create more interesting strings, brass, and woodwinds but it also becomes possible to create vocals, inharmonic sounds such as bells, and countless other acoustic and synthetic sounds. One simple fact is that a single oscillator will have harmonics where each successive partial has less amplitude. While the higher harmonics of acoustic instruments do diminish in amplitude they do not follow this progression so rigidly. Below are harmonic diagrams of some acoustic instruments.

Synthesis Through Harmonic Analysis and Reverse Engineering

The use of harmonic analysis in synthesis is nothing new, yet it has not been used much outside of academic research and professional patch programming. There are a few reasons for this. The first is that harmonic analysis requires a fair amount of processing power that until recently was too expensive for most of us. Fortunately home computers are now powerful enough to perform the analysis in real-time while using only a small percentage of the computer's processor. The second reason is that it is generally perceived as difficult to use. There is quite a bit of mathematical theory that goes into computing the Fast-Fourier Transforms (FFT) on which harmonic analysis is based but as far as using harmonic analysis as a tool for synthesis it's rather straightforward. The third reason, somewhat as a result of the first two, is that it simply hasn't become part of convention.

This is super powerful stuff and it has the potential to change the way you program sound. Even though this section is centered on analyzing harmonics and reverse engineering musical content it should be noted that it will also improve your ability to create original patches as well. After you use it for a while and you come to understand the spectral makeup of different sounds you will rely less and less on having to look at harmonic diagrams. They'll be in your head. This book is about analog synth programming but it should be mentioned that harmonic analysis also makes FM synth programming infinitely easier to understand. Don't even bother programming FM without it! I digress…

Before proceeding I recommend that you first get a hold of software that can perform harmonic analysis/FFT in realtime and play around with it for a bit. Most of the more sophisticated sound editors provide the ability to display harmonic diagrams of sound files albeit usually not in real-time. For real-time analysis I highly recommend the Fre(a)koscope VST plug-in. This plugin was used to create all of the harmonic diagrams in this book. It has an excellent interface, is easy to use, and it's free! If you don't have a host that supports VST plugins be sure to visit kvraudio.com where you can find a number of hosts available for free. If you have any questions please email me at synthcookbook@yahoo.com

The analysis/synthesis process can be broken down into three stages for each sound.

1. The most important step. Obtain a harmonic diagram of the sound that is to be emulated or recreated. Pick the best oscillator waveforms, pitches, and amplitudes to match the diagram.
2. High frequency harmonics are generally the last to rise in amplitude and are the first to die back out. While comparing with the harmonic diagram of the original source use the synth's filter envelope to match this behavior.
3. Using a sound editor, match the synth's envelopes to the original source.

Step one is essential while steps 2 and 3 are not always possible or required. Envelopes may not require any visual analysis since they can be duplicated by ear with good results more easily than harmonics. Note: The lowest-pitched harmonics are the most important. As a general rule of thumb try to get the first 5-10 harmonics as close to their corresponding levels in the original source.

Reverse engineering a patch from another synth:

The easiest sounds to recreate are those from other analog synths. Because analogs have very similar feature sets it makes for higher accuracy. About ten years ago I had to sell a Sequential Circuits Six-Trak but I made sure to record samples of some of my favorite patches that I'd programmed for it. My favorite was a lead with thick unison and we'll begin by reverse engineering this patch. The sample of this patch can be found on the CD.

Playing the sample through Fre(a)koscope produces the following harmonic diagram,

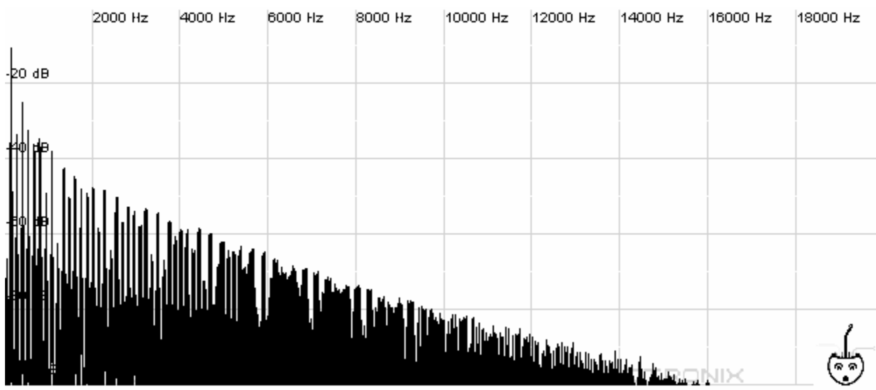

Figure 37. Harmonics of a patch originally programmed on a Sequential Circuits Six-Trak

Here Fre(a)koscope has been set to a linear display and a maximum window size with a frequency range of 20 Hz – 20 kHz. The lower harmonics are the most important so let's decrease the maximum frequency of the display so that we can see these in greater detail. Lower the maximum range down to about 6 kHz.

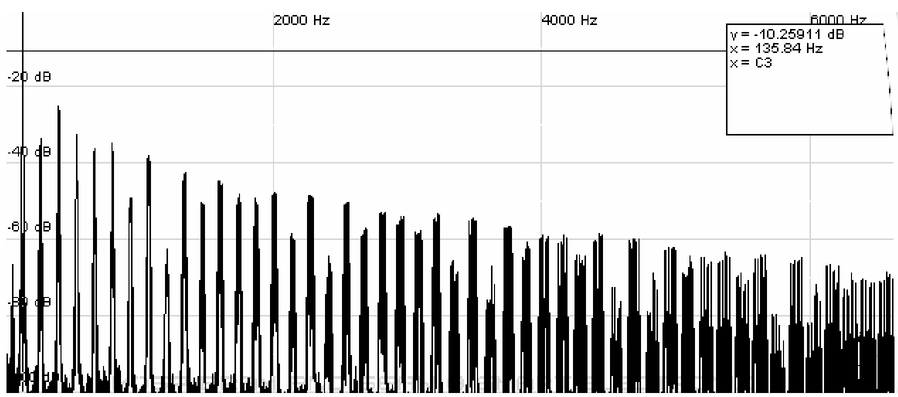

Figure 38. Six-Trak patch zoomed in to show finer detail in the lower-frequency harmonics. The pitch of the sampled note is at C3 on the keyboard

In Fre(a)koscope it is possible to get decibel value, frequency, and equivalent key number by clicking on a section of the graph. In the graph above the fundamental has been selected and from the inset box at the upper right we see that the sample is tuned to the C3 key. We need to know this so that we know what key to play on the synth that is being programmed. There are many important details in the above graph that give us a clue as to where we should start. Notice that both even and odd harmonics are present. This rules out using a single triangle wave or square wave since they have only odd harmonics.

The two remaining options are sawtooths and pulses since they possess both even and odd harmonics.   Also notice that the higher harmonics become broad and are not the narrow spikes that we would expect with pure sawtooths or pulses. This indicates that there could be some pulse-width modulation taking place or the oscillators have been set to unison. We will get back to the broad harmonics in a bit but for now look at the way the harmonic amplitudes do not decrease evenly from left to right. A sawtooth that has been synced can produce harmonics that look like this but so can a pulse with variable width. Let's take a look at the synced sawtooth option first. Set your synth up with the slave oscillator synced to the master oscillator and have the slave produce a sawtooth wave while at the same time the master's output is turned all the way down. Increase the pitch of the slave while keeping an eye on the synth's harmonic analysis. The closest that we can get to the original harmonics this way is when the slave oscillator is tuned about 8 semitones above the master.

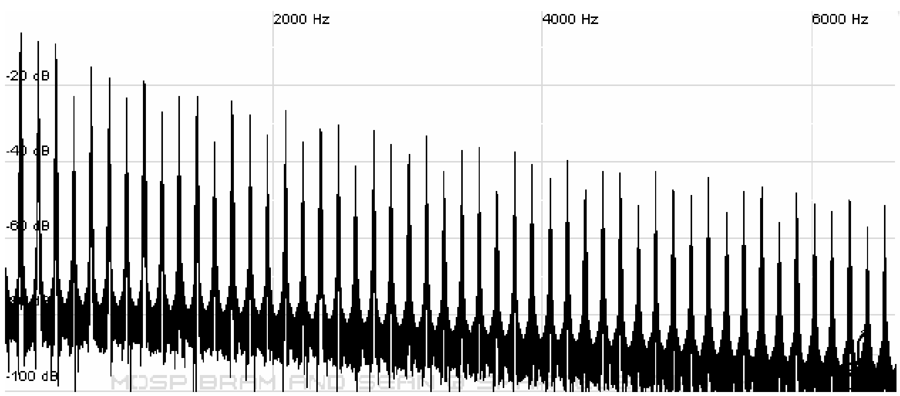

Figure 39. A synced sawtooth oscillator tuned 8 semitones above the master produces results similar to the original patch but not close enough

What are the notable differences between this graph and that of the original Six-Trak sample? Of course the harmonics are all narrow and again we will get to that in a moment but for now we only need to compare their amplitudes to the original patch. The higher harmonics in our patch are slightly louder than those in the original. Notice that they are higher on the graph. This can be remedied by applying some low-pass filtering but there is a bigger discrepancy than this. Notice that the harmonics of our patch do not undulate up and down in quite the same fashion as the original patch. It would however be a good guess that if we find out what is broadening the higher harmonics of the original patch and we apply it to our patch as it is right now that the two would nonetheless sound very similar but we are going to try to do better. Remember this was originally produced by another synthesizer so we should be able to get really close and we can probably even nail it.

Let's see what we can come up with using a pulse instead of a sawtooth. Again only use one oscillator but this time do not sync it to the other. Set it up to produce a pulse. Begin by setting the pulse width to produce a square wave. Remember that a square wave should only show odd harmonics. Now very slowly begin decreasing the pulse width. Before the width has been decreased much at all you should get harmonics that are identical in amplitude to those of the original patch. This happens to be at a pulse width of about 44%, something not much narrower than a square wave but harmonically different. Now we haven't yet applied any filtering so the harmonics will not diminish quite as rapidly as those in the original patch but the levels of each harmonic relative to their neighbors follows the same pattern as those in the original.

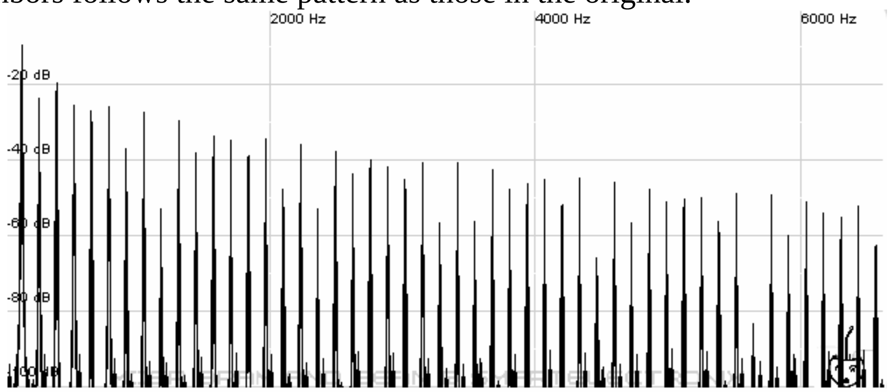

Figure 40. A pulse width of 44% produces harmonics that are very similar to the original patch

Next experiment with some low-pass filtering until the level of the higher harmonics matches those of the original. A 24 dB filter set to a cutoff frequency of 2.5 kHz produces the following graph.

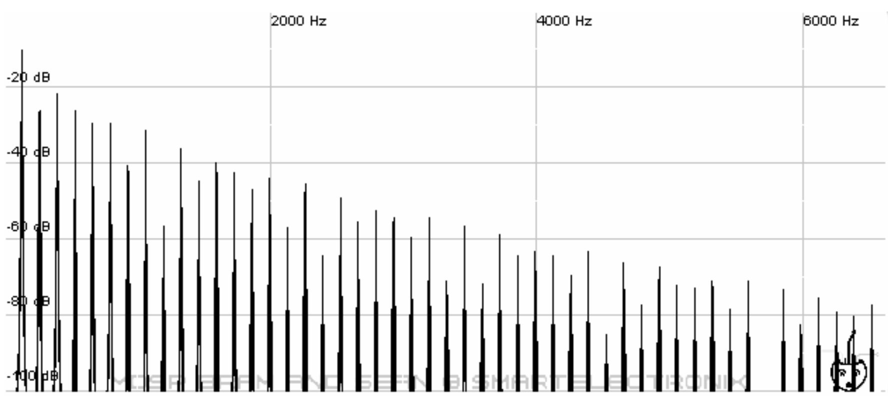

Figure 41. Apply some low-pass filtering to match the harmonic amplitudes to those in the Six-Trak patch

There are some differences between this and the original but the amplitudes are pretty darn close. While two synths can produce *practically* the same results they cannot produce *exactly* the same results.

The one last thing that we need to take care of is determining what is causing the higher harmonics to broaden. This could either be due to pulse-width modulation or oscillator

HOW TO PROGRAM THE COOKBOOK PATCHES

Oscillators

Wave : Waveform choices include sawtooth, square, triangle, and pulse width (PW). Pulse width values are given as a percentage of the waveform period. For instance a pulse width of 50% is a square wave and a pulse width of 25% is a pulse with a width that is half that of a square and sounds more "brassy." Pulse widths greater than 50% are not included as they are just inverted versions of those between 0% and 50% and therefore sound the same. You will need to know if your synth's pulse-width setting goes from 0%-50% or, more commonly, 0%-100%. If you have a synthesizer such as the ImpOscar which uses values of 50%-100% simply subtract the pulse width values in the book from 100 to get the value you need to use. If the synthesizer lacks the ability to set the pulse width of the square wave then substitute pulse-width values 0%-30% with a sawtooth wave and values from 30% - 50% with a square wave. If the synthesizer lacks a triangle wave then use a square wave in it's place when called for along with slightly more low-pass filtering. Triangle and square waves sound similar because they are both composed of odd harmonics but the higher harmonics of the triangle are less audible.

Tune: Values are listed in octaves, semitones, and cents. This is the value by which one oscillator should be detuned from the other oscillator. For patches where both oscillators have a tuning value listed it is the amount by which each oscillator is to be detuned from a common pitch. If the tracking for oscillator 2 is turned off this field will list the specific pitch for which oscillator 2 is to be tuned.

Mix: This is the amplitude of the oscillator. Values are given in decibels where the maximum amplitude is 0 dB and quieter amplitudes are given as negative values. Some synthesizers display amplitudes where the maximum amplitude is not 0 dB. It may for instance be some positive value. In these cases just subtract the value listed in this book from the maximum value displayed by the synthesizer. For example if the synthesizer displays a maximum value of 13 dB for the volume of one of the oscillators and the setting called for in this book is -7 dB then the oscillator should be set to a value of 6 dB. It's not that big of a deal just as long as you make sure the oscillators are separated by the same amount. The percentage to the right of the decibel value is the percentage of the maximum setting of the oscillator amplitude. This can be used for synthesizers that do not display values in decibels. Since oscillators on different synths have amplitude adjusters with different responses this percentage is a ballpark approximation and may require some adjustment by ear. For more accuracy it is best to calibrate the synths mix settings as explained in the last section of this book.

Oscillator 2 tracking: This is either on or off. When on the pitch of oscillator 2 will track the keyboard. When off oscillator 2 will play the same pitch on all keys.

Oscillator 2 sync: When on it syncs oscillator 2 to oscillator 1. You may need to check your synthesizer's documentation to be sure which of its oscillators syncs to the other when syncing is turned on.

Noise

Noise: White noise source set to either on or off.

Mix: Amplitude of the noise source given in decibels. The statements made above in the section on oscillator mix apply here as well.

LFO

Routing: Sets the destination of the LFO. Destinations include amplitude, pitch, filter cutoff, and pulse width.

Wave: Waveform used by the LFO as a modulation source. Waves used are sine, triangle, and square.

Frequency: Sets the rate of the LFO. Values are given in Hertz (Hz). Qualitative suggestions are given in italics.

Depth: Sets the amount by which the destination is modulated by the LFO. This is given as a percent of the settings range when modulating amplitude, filter cutoff, and pulse width. The knob/slider that controls depth would be turned half way for a depth of 50%. When pitch is set as the destination the depth is given in terms of the amount of pitch that a played note will vary.

Glide: Also known as portamento. Set to either on or off

Time: Glide time. The amount of time that it takes for the pitch to change from one struck note to the next.

Unison: Set to on or off. Your synthesizer may allow you to control the number of oscillators that are used. Usually 4 to 6 oscillators sound good.

Voices: The polyphony of the patch. Either set to a specific number, mono, or multi. Multi means there is no restriction on the amount of polyphony.

Low-Pass Filter

Cutoff: Values are given for 24 dB filters and 12 dB filters. Values are given in hertz and kilohertz and go from 20 Hz to 20 KHz. If your synthesizer uses something other than a 24 dB or 12 dB filter you can figure out the required setting by extrapolating from the given values for 24 and 12 dB filters. Generally, the smaller the amount of attenuation (the decibel number) the lower the cutoff frequency should be. With a greater amount of attenuation set the cutoff to a higher frequency. For example, if the listing for a 12 dB filter calls for a cutoff of 200 Hz then a 6 dB filter would be set to a frequency less than 200 Hz. A 36 dB filter would be set to a cutoff frequency higher than what is listed for the 24 dB filter value, etc.

To the right of each cutoff frequency is listed the percentage of the full dial position. Of course each synthesizer's filter cutoff will follow a different response curve so this value is provided as a rough estimate. It is based on the average taken from ten different synthesizers. Actual values can vary drastically. ***Following these percentages will require at least some adjustment by ear***. If you patch a sound and it doesn't sound right then most likely it is due to a discrepancy between the filter-cutoff setting percentages and the actual filter cutoff of the synth. As a simple test I would recommend patching the Angels patch on pg 33. If it sounds more like a brass instrument than a vocal then the listed percentages are higher than what your synth produces. Set the cutoff for each patch below the listed percentage value. If the angels sound like they are being smothered then the listed percentages are too low. Set the cutoff for each patch above the listed percentage value. To achieve higher accuracy on synthesizers that do not display cutoff in hertz it is recommended to follow the calibration instructions at the end of the book. This calibration procedure for the filter while not critical is HIGHLY recommended! If you only calibrate one parameter it should be the filter cutoff.

Resonance: The amount of resonance applied to the filter. Values are given as a percentage where 0% applies no resonance and 100% applies the maximum amount.

Envelope: Sets the amount by which the filter envelope controls the amount of filtering. On some synths this is called the Initial Level setting. Again 0% is the minimum value and 100% is the maximum.

ADSR: Attack, Decay, and Release are given in terms of time with values in seconds. Sustain is given as a percentage of its maximum amount.

Amplifier

ADSR: The same as the envelope for the filter. Attack, Decay, and Release are given in terms of time with values in seconds. Sustain is given as a percentage of its maximum amount.

Alternate Patch Settings

Some of the patches in the book have alternate patch settings which are intended to be used if the synth does not have the feature listed in the primary patch. For instance let's say that you are programming the cello patch which calls for oscillator 1 to use a pulse with a width of 10% but your synth doesn't have the ability to produce a pulse. Use the alternate setting which calls for a sawtooth wave instead.

Note about times: Some synthesizers list times in seconds (s) and others list times in milliseconds (ms). This book uses seconds but the conversion is straightforward. Simply move the decimal three places. The relation is:

1 s = 1000 ms

0.1 s= 100 ms

0.01 s = 10ms

0.001 s = 1 ms

So for example 0.08 s is the same thing as 80 ms, 0.28 s is also 280 ms, and 1.35 s is 1350 ms, etc.

STRINGS

Banjo

Oscillator 1		LFO		Low-pass filter			
Wave:	PW = 20%	Routing:	Amplitude		24db	12db	
Tune:	-	Wave:	Triangle	Cutoff:	2.9khz *72%*	1.5khz *63%*	
		Frequency:	10hz *fast*	Resonance:	0%		
Mix:	-13.9db *60%*	Depth:	10%	Envelope:	75%		
Oscillator 2		Glide:	-	**A**	**D**	**S**	**R**
Wave:	PW = 10%	time:	-	0s	0.19s	0%	0.19s
Tune:	+5 semi	Unison:	off	**Amplifier**			
		Voices:	5	**A**	**D**	**S**	**R**
Mix:	-21.5db *45%*			0s	0.67s	0%	0.67s
Osc 2 Track:	on						
Osc 2 Sync:	**on**						
Noise:	off						
Mix:	-						

Cello

Oscillator 1		LFO		Low-pass filter			
Wave:	PW = 10%	Routing:	Amplitude		24db	12db	
Tune:	-	Wave:	Sine	Cutoff:	40hz *10%*	40hz *10%*	
		Frequency:	7.5hz *moderate*	Resonance:	0%		
Mix:	0db *100%*	Depth:	5%	Envelope:	90%		
Oscillator 2		Glide:	off	**A**	**D**	**S**	**R**
Wave:	Square	time:	-	0s	3.29s	78%	max
Tune:	-	Unison:	off	**Amplifier**			
		Voices:	multi	**A**	**D**	**S**	**R**
Mix:	0db *100%*			0.06s	max	100%	0.30s
Osc 2 Track:	on						
Osc 2 Sync:	off						
Noise:	off						
Mix:	-						

Alternate: Osc 1: sawtooth

BRASS

French Horn

Oscillator 1		LFO		Low-pass filter			
Wave:	PW = 10%	Routing:	-		24db	12db	
Tune:	-	Wave:	-	Cutoff:	40hz *10%*	40hz *10%*	
		Frequency:	-	Resonance:	20%		
Mix:	0db *100%*	Depth:	-	Envelope:	45%		
Oscillator 2		Glide:	off	**A**	**D**	**S**	**R**
Wave:	-	time:	-	0.05s	5.76s	94%	0.39s
Tune:	-	Unison:	off	**Amplifier**			
		Voices:	multi	**A**	**D**	**S**	**R**
Mix:	-			0s	3.9s	96%	0.93s
Osc 2 Track:	-						
Osc 2 Sync:	-						
Noise:	off						
Mix:	-						

Alternate: Osc 1: sawtooth, Osc 2: triangle, +2 octaves, Fiter: envelope 55%

Harmonica

Oscillator 1		LFO		Low-pass filter			
Wave:	PW=2%	Routing:	PW osc 1		24db	12db	
Tune:	-	Wave:	Triangle	Cutoff:	1.9khz *66%*	1.2khz *59%*	
		Frequency:	1.6hz *slow*	Resonance:	0%		
Mix:	-29db *30%*	Depth:	85% *deep*	Envelope:	65%		
Oscillator 2		Glide:	off	**A**	**D**	**S**	**R**
Wave:	PW=15%	time:	-	0.16s	max	100%	0.16s
Tune:	+8 semi	Unison:	off	**Amplifier**			
	+7 cents	Voices:	multi	**A**	**D**	**S**	**R**
Mix:	-18db *55%*			0.13s	0.33s	50%	0.14s
Osc 2 Track:	on						
Osc 2 Sync:	**on**						
Noise:	off						
Mix:	-						

CALIBRATION

This section is intended for those synthesizers that do not give the values of their settings in decibels, hertz, or seconds. By first calibrating such a synthesizer the patches in this book can be programmed much more accurately than using the suggested percentage settings that are given in italics. It is advisable to at the very least calibrate the oscillator mix and filter cutoff frequency. There are two parts to this section. The first part explains how to calibrate a synth using the sound files on the CD and the second part explains how to using a sound editor and other metering devices. Using the CD is probably the quickest way to get started while using the directions in the second part will produce slightly more accurate results.

The most optimal way to use these two sections is probably to use the CD to calibrate the filter's cutoff frequency and if you can record your synth into a sound editor then use Part B to determine oscillator mix levels, LFO frequency and ADSR values. If you would like more clarification on any part of this section please feel free to email me at *fred_v_welsh@yahoo.com*

Part A - Using the CD

Using the sounds files on the CD is quite simple. In a nutshell you are basically just comparing the output of your synth to each file and figuring out values such as decibels, Hz, and seconds from the filename of the sound that most closely matches what you hear from your synth.

Oscillator Mix:

A. Get a piece of paper and a pencil or pen.

B. Go into the *Oscillator Mix* folder on the CD. Although the filter isn't being calibrated at this point it is still important to use samples that correspond to the synth's filter type. Open the folder that corresponds to the filter type used by your synth. If both filter types are available then you can just pick one but of course make sure the synth is set to that filter.

C. Now you should see two folders. We are going to use the *Dual Sawtooth Waveforms* folder. The *Dual Sawtooth Waveforms* folder contains samples where both oscillators are set to sawtooth and one is raised two octaves. Set your synthesizer up with both oscillators set to sawtooth, one raised two octaves. The purpose of this is so that the amplitude of one oscillator can be compared to the other

D. In each one of these samples the high-pitch oscillator is left at its maximum setting. This maximum setting corresponds to 0 dB. The name of each file indicates the amplitude of the low-pitch oscillator. On your synth, turn the amplitude of both sawtooths all the way up. It should now produce the same tonal characteristics as the 0 dB file when you play middle C on the keyboard. Also make sure your filter is completely open and that amplitude velocity sensitivity is off.

E. Write down the number, dial position, or whatever your synth uses that corresponds to the maximum setting and next to it write down "0 dB." Make sure to write neatly because you are going to be using these values to patch the sounds in this book.

F. Listen to the tonal characteristics of the -5 dB file. Now lower the amplitude of your low-pitch oscillator until it produces the same tonal characteristics as the -5 dB file. Write down the position followed by "-5 dB."

G. Continue to do the same thing for each of the remaining sound files

The "*Single Sawtooth Waveform*" folder contains samples of a single sawtooth oscillator which just like in the instructions above is lowered in 5 dB increments. This is just an alternative to the Dual Sawtooth instructions above and is carried out in the exact same way.

Filter

Determining Filter Type

If you already know whether your filter is of the 12 dB or 24 dB variety then you can skip this section and proceed to the *Determining Filter Cutoff* section.

A. Open the "*Filter*" folder and then open the "*Determine Filter Type*" folder. There are four files in here. For each set (12 dB or 24 dB) there is a file where the filter resonance is set to a pitch of 440 Hz and after about ten seconds the resonance is turned all the way down. There are also two files where the filter is completely open at 20,000 Hz and resonance is set to zero. Play both of these files and you'll notice that the 12 dB filter produces a slightly brighter sound. That's because a filter that is "open" isn't completely open! It will still have some effect on the tone.