NUTS and BOLTS PART 16

Pitch Shifting

BY ALEX CASE

We all know what happens when you play an audio tape at a faster speed than it was recorded: the pitch of the recording goes up. Play it slower than the recorded speed and the pitch goes down. Somewhere in this simple principle lies an opportunity for audio exploration and entertainment.

Can it be done digitally? Of course. How is it done? Using a delay. Yes, I said delay. Loyal readers of Nuts & Bolts have just spent the last three issues reading about delays: echo, chorus, flange, and their many cousins. Yet we aren't finished discussing delay, because even pitch shifters are built on this effect.

To see how, you've got to put up with a bit of math (which I seem to sneak into every article) and follow along with Figure 1. Figure 1i shows a simple sine wave with a chosen frequency of 250 Hertz. This sine wave completes a cycle every 4 milliseconds, confirmed courtesy of the following familiar equation:

Period = the time to complete one cycle

Period = 1 / Frequency

Period = 1 / 250 Hz

Period = 0.004 seconds or 4 milliseconds

Figure 1i labels some key landmarks during the course of a single cycle. We start the clock at point A.

Here, at time equals zero, the sine wave is at an amplitude of exactly zero and is increasing. It reaches its positive peak amplitude at B, taking exactly 1 millisecond to do so. It has an amplitude of zero again at the halfway point (time equals 2 milliseconds) labeled C. This makes it look a lot like point A, but while the amplitude is the same, notably zero, it is decreasing this time.

D is the point of maximum negative amplitude and occurs 3 milliseconds after the beginning of the cycle. And E has our sine wave returning 4 milliseconds after point A to what looks exactly like the starting point: the amplitude is zero, and increasing. Okay so far? We are going to follow these points through some signal processing and move them around a bit.

Fixed delay

Run this sine wave through a fixed delay of, say, 1 millisecond, and you get the situation described in Figure 1ii. Visually, you might say the sine wave slips along the horizontal time axis by 1 millisecond.

Looking point by point, Table 1 shows us what happens. Point A originally occurred at a time of zero. Introduce a fixed delay of one millisecond and Point A now occurs at time equals one millisecond.

The other points follow. Point D, for example, occurred undelayed at a time of 3 milliseconds. After it has been run through a fixed, unchanging delay of one millisecond, Point D is forced to occur at a time of 4 milliseconds.

Accelerating delay

Here's the mind bender. What happens if the delay isn't fixed? What if the delay sweeps from a starting time of 1 millisecond and then increases, and increases, and increases...?

Table 2 summarizes. Here the delay changes at a rate of one millisecond per millisecond. Say what? For every millisecond that ticks by during the course of this experiment, the delay gets longer by one millisecond. For example, if at one point the delay is 10 milliseconds, then five milliseconds later the delay unit is operating at 15 milliseconds.

If you haven't had the chance to study physics, you might be puzzled by the idea of changing the delay time at a rate of 1 millisecond per millisecond. I find it helpful here to get in my car.

Say you are driving at a speed of ~~85~~ (edited for your safety) 55 miles per hour and accelerating at a rate of 1 mph per hour (our Canadian neighbors should use kilometers per hour for similar results). That means that for each hour that passes by your speed increases by one mile per hour. Driving 55 mph now becomes 56 mph an hour from now. Whoa!

If you subscribe to the idea that cops won't pull you over for speeding until you are at least 10 miles per hour over the speed limit, then—starting at the speed limit—you can drive most of the way across Texas (it's about 600 miles from Dallas to El Paso) without getting a ticket.

Back to the music. Here we are increasing the delay by 1 millisecond each millisecond. And the surprising result is a change in the pitch of the track.

Table 2 shows the location of our sine wave landmarks both before and after the introduction of this steadily increasing delay. Point A initially occurs at a time of zero. At this time the delay is also zero. Point A then remains unchanged and occurs at time zero.

Skip to point C. It originally occurs at a time of two milliseconds. By this time the delay has increased from zero to two milliseconds. This delay of two milliseconds leads point C to finally occur at a time of four milliseconds after the beginning of this experiment.

Do the math point by point and you get a sine wave that looks like Figure 1iii. The key landmarks are identified. The result is clearly still a sine wave. But since it takes longer to complete the cycle, we know the pitch has changed.

Back to our trusty equations. Looking at the new sine wave at the bottom of Figure 1, let's calculate its frequency. The sine wave in Figure 1iii takes a full 8 milliseconds to complete its cycle.

Period = the time to complete one cycle
Period = 8 milliseconds or 0.008 seconds
Frequency = 1 / Period
Frequency = 1 / 0.008
Frequency = 125 Hertz

That's right. The constantly increasing delay caused the pitch of the signal to change. We sum it up as follows. A 250 Hertz sine wave run through a delay that increases constantly at the rate of 1 millisecond per millisecond will be lowered in pitch to 125 Hertz. Strange, but true.

Skipping the details—though you are encouraged to prove these on your own—we find a number of finer points on pitch shifting. Our example demonstrated that an increasing delay lowers the pitch. It is also true that a decreasing delay raises the pitch.

In addition, our example found that changing the delay at a rate of 1 millisecond per millisecond moved the pitch by an octave. It is also possible to change the pitch by two octaves, or a minor third, or a perfect fifth—whatever you desire; one need only change the delay at the correct rate.

So the underlying methodology of pitch shifters is revealed. A pitch shifter is a device that changes a delay in specific controlled ways so as to allow the user to affect the pitch of the audio.

Naturally, it ain't easy. We've got a key problem. Return to our example in which we lowered the 250 Hz sine wave by an octave through a steadily increasing delay. If we imagine applying this effect to an entire three and a half minute tune, not just a single cycle of a sine wave, we find we are increasing the delay from a starting point of one millisecond to a final delay time of 210,000 milliseconds (3$^{1}/_{2}$ minutes equals 210,000 milliseconds).

That is, at the start of the tune we add an increasing delay: 1 ms, then 2 ms, and so on. By the end of the tune, we are adding a delay of 210,000 milliseconds. This highlights two problems.

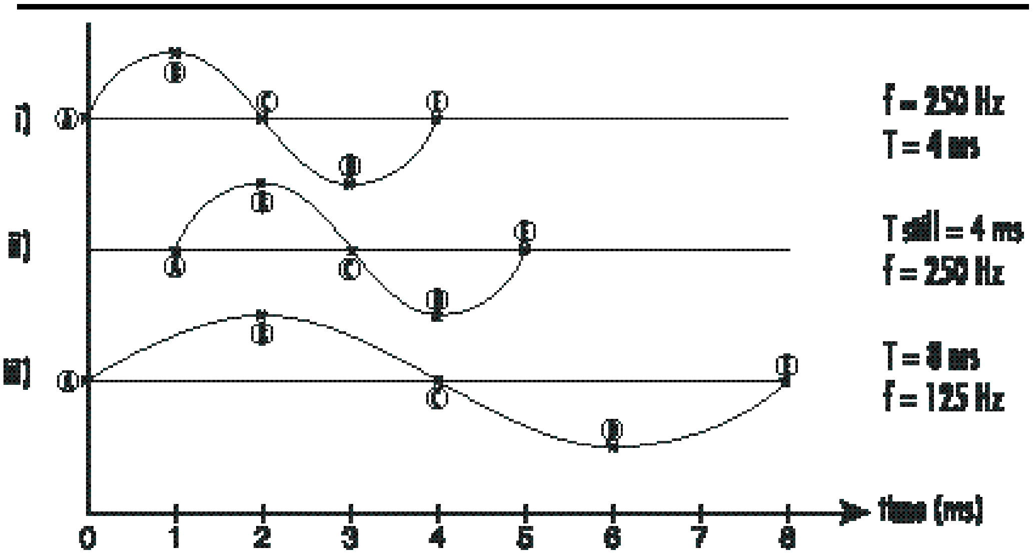

FIGURE 1: PITCH SHIFTING THROUGH A VARIABLE DELAY

First, we need a very long delay. Most delays are capable of a one second delay (1,000 milliseconds) at the most. Super cool mega turbo delays might go up to maybe 10 seconds of delay. But a delay of hundreds of thousands of milliseconds (hundreds of seconds) is a lot of signal processing horsepower that is rarely available—RAM isn't that cheap.

Second, our song, which used to be 3$^{1}/_{2}$ minutes long, doubles in length to seven minutes as we lower the pitch by one octave. Consider the last sound at the very end of the song.

Before pitch shifting it occurred 3$^{1}/_{2}$ minutes (210,000 milliseconds) after the beginning of the song. By this time our pitch shifting delay has increased from 1 millisecond to 210,000 milliseconds. Therefore the final sound of the pitch shifted song occurs at 210,000 milliseconds (original time) plus 210,000 milliseconds (the length of the delay). That is, the song now ends 420,000 milliseconds (that's seven minutes!) after it began. The 3$^{1}/_{2}$ minute song is lowered an octave but doubled in length.

Simply increasing the delay forever as above is exactly like playing a tape back at half the speed it was

NUTS and BOLTS

that increases at rate of 1 millisecond per millisecond will lower the audio by one octave. Any delay time. So why not keep it a small delay time?

The devil is in the details. Getting the pitch shifter to reset itself in this way without being noticeable to the listener isn't easy. It is a problem solved by clever software engineers who find ways to make this inaudible.

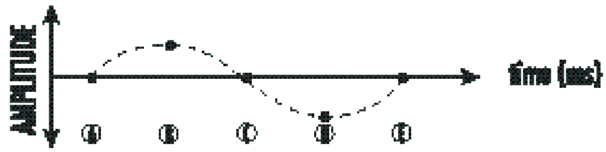

Undelayed 250 Hz Sine Wave	0	1	2	3	4	(Figure 1i)
+ Fixed 1 ms Delay	1	1	1	1	1	(Figure 1ii)
250 Hz Sine Wave with a Fixed 1 ms Delay	1	2	3	4	6	

TABLE 1: THE TIME WHEN OUR SINE WAVE LAND MARKS OCCUR, WITH AND WITHOUT A FIXED DELAY (FREQUENCY = 250 Hz, PERIOD = 4 ms DELAY FIXED AT 1 ms)

recorded. The pitch goes down and the song gets longer.

Where pitch shifting signal processors differentiate themselves from tape speed tricks is in their cleverness here. Digital delays can be manipulated to always increase, but also to reset themselves.

In our sine wave example, what happens if our digital delay increases at a rate of exactly one millisecond per millisecond but never goes over 50 milliseconds in total delay? That is, every time the delay reaches 50 milliseconds it resets itself to a delay of zero and continues increasing from this new delay time at the same rate of one millisecond per millisecond.

The result is pitch shifting that never uses too much delay and never makes the song more than 50 milliseconds longer that the un-pitch shifted version. After all, our analysis showed it was the rate of change of the delay that led to pitch shifting, not the absolute delay time itself. Any delay time

Older pitch shifters 'glitched' as they tried to return to the original delay time. These days—we are lucky to be alive in audio in the year 2000—those glitches are mostly gone. Today we simply reach for a device labeled "pitch shifter," dial in the desired settings (gimme a shift up of a major third mixed in with a shift down of a perfect fourth), and get to work. Life is good.

Side effects

Before we get into the special effects we create with pitch shifting devices, it is worth noticing that pitch shifting is a natural part of some effects we've already investigated here in the Nuts & Bolts series. Recall the chorus effect that comes from adding a slowly modulated delay of about 30 milliseconds.

As you listen to the richness that the chorus effect adds to a vocal or guitar, listen for a subtle amount of pitch shifting. That's right, pitch shifting is a component of that effect we call chorus.

Since a chorus pedal relies on a modulating delay, it introduces a small amount of pitch shifting. As the delay time sweeps up, the pitch is slightly lowered. As the delay time is then swept down, the pitch is then raised, ever so slightly. Repeat until thickened.

Special effects

The Nuts & Bolts review of the a basic pop mix, 'Mixing by Numbers' in the 4/00 issue, introduced a common effect built in part on pitch shifting: the spreader.

A quick review of this effect: the spreader is a "patch" that enables you to take a mono signal and make it a little more stereo-like. You 'spread' a single track out by sending it through two delays and two pitch shifters.

The delays are kept short, each set to different values somewhere between about 15 to 50 milliseconds. Too short and the effect becomes a flange/comb filter (as we discussed last month). Too long and the delays stick out as distinct echoes. So our window for acceptable delay times in this effect is between about 15 and 50 milliseconds.

In using a spreader, the return of one delay output is panned left while the other is panned right. The idea is that these quick delays add a dose of support to the original monophonic track. In effect, these two short delays simulate some early sound reflections that we would hear if we played the sound in a real room. The spreader takes a single mono sound and sends it to two slightly different short delays to simulate reflections coming from the left and right.

That's only half the story. The effect is taken to the next level courtesy of some pitch shifting. Shift each of the delayed signals ever so slightly, and the formerly boring mono signal becomes a much more immersive, interesting loudspeaker creation.

Detune each delay by a nearly imperceptible amount, maybe 5 to 15 cents. The goal of the spreader is to create a stereo sort of effect. As a result, we try to keep the signal processing on the left and right sides ever so slightly different from each other. Just as we dialed in unique delay times for each side of this effect, we dial in different pitch shift amounts as well—maybe the left side goes up 8 cents while the right side goes down 8 cents.

Like so much of what we do in recording and mixing pop music, the effect has no basis in reality. By adding delay and pitch shifting, we aren't just simulating early reflections from room surfaces anymore. The spreader makes use of our signal processing equipment (delay and pitch shifting) to create a big stereo sound that only exists in loudspeaker music. This sort of thing doesn't happen in symphony halls, opera houses, stadiums or pubs. It's a studio creation, plain and simple.

Take this effect further and you end up with what I think of as a 'thickener.' Why limit the patch to two delays and two pitch changes. What if you have the signal processing horsepower in your DAW or in your racks of gear to chain together eight or more delays and pitch shifts?

Try it. While it'll sound unnatural when used on vocals, many keyboard parts respond well to the thickening treatment. Modulate those delays like a chorus and, guess what? More pitch shifting is introduced.Added in small, careful doses, this densely packed signal of supportive, slightly out-of-tune delays will strengthen and widen the loudspeaker illusion of the track.

Big time

Enough with these subtle pitch changes. Let's add a serious amount of pitch shifting.

Hammond B3 organs and many blues guitars are often sent through a rather wacky device: the Leslie Cabinet. The Leslie is a hybrid effect that is built on pitch shifting, volume fluctuation, and often a good dose of tube overdrive distortion.

It is essentially a tricked out guitar or keyboard amp in which the speakers rotate. Honest. The high frequency driver of a Leslie is horn loaded, and the horn spins around within the amp.

Crazy as it sounds, the engineers who came up with this were really thinking. It would be very difficult to spin the large woofer to continue the effect at low frequencies. Instead they enclosed the woofer inside a drum. The drum has holes in it and rotates. The result is a low frequency simulation of what the Leslie is doing with the horns at higher frequencies, er, well, sort of.

The Leslie is too funky a device to cover in detail now, but we mention it because it is part of our pitch shifting toolkit.

So what's it sound like? With the drum and horn rotating, the loudness of the music increases and decreases at any given listener position—amplitude modulation. And with the high frequency horn spinning by, a Doppler effect is created: the pitch increases as the horn comes toward the listener/microphone and then decreases as the horn travels away.

The typical example used in the study of the Doppler effect is a train going by, horns ablaze. That classic sound of the pitch dropping as the train passes is based on this principle. Sound sources approaching with any appreciable velocity will increase the perceived pitch of the sound. As the sound source departs, the pitch similarly decreases. The net result of the Leslie system then is a unique fluttery and wobbly sound.

The Leslie effect is used wherever B3s and their ilk are used. Typically offering two speeds of rotation, you can hear a fast Leslie and a slow Leslie effect, as well as the acceleration or deceleration in between.

Listen to the single note organ line at the introduction to 'Time and Time Again' on the Counting Crows' first record, August And Everything AfterThe high note enters with a fast rotating Leslie. As the line descends, the speed is reduced. Listen carefully throughout this song, this album, and other

NUTS and BOLTS

B3-centric tunes and you'll hear the Leslie pitch shifting vocabulary that keyboardists love.

Of course, you can apply it to any track you like—guitar, vocal, oboe—if you have the device or one of it many imitators or simulators.

The hazard with an obvious pitch shift is that it can be hard to get away with musically. You've heard special effects in the movies and on some records, where a vocal is shifted up or down by an octave or more. Too low and it conjures up images of death robots invading the mix to eat Shreveport. Too high and your singer becomes a gremlin-on-helium sort of disaster.

pitch shifter makes it easy to add a constant,never ending, never changing, just plain annoying harmony line.

Dial in a pitch change that is a major third up and add it to the lead vocal. If the song is entirely diatonic within a major key and is a very happy song, I mean very happy, this might work. Otherwise it is going to sound cloyingly sweet, like adding maple syrup to the ice cream you put on top of your shoo fly pie.

The trick to creating harmonies using pitch shifting is to compose musical harmonies. And a static pitch shift will rarely cut it.

Fortunately, devices and software plug-ins to facilitate this abound (probably the most famous is the DigiTech Vocalist series). The pitch shifting can essentially be tied to MIDI note commands enabling you to dictate the harmonies from your MIDI controller. The pitch shifter is processing the vocal line on tape or disk according to the notes you play on the keyboard.

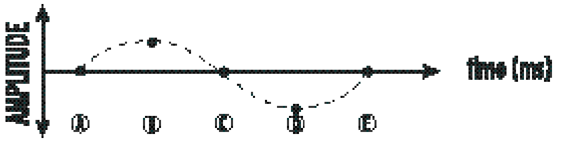

	A	B	C	D	E	
Undelayed 250 Hz Sine Wave	0	1	2	3	4	(Figure 1i)
+ Increasing Delay 1 ms/ms)	0	1	2	3	4	(Figure 1ii)
125 Hz Sine Wave	0	2	4	6	8	

TABLE 2: THE TIME WHEN OUR SINE WAVE LANDMARKS OCCUR, WITH THE ADDITION OF A DELAY THAT INCREASES BY A MILLISECOND FOR EACH MILLISECOND THAT TICKS BY

In the hands of talented musicians, aggressive pitch shifting really works. TANKAPA (The Artist Now Known As Prince Again) lowers the pitch of the lead vocal track and takes on an entirely new persona in the song, Bob George from The Black Album. The effect is obvious. The result is fantastic.

No effort was made to hide the effect in the bass line of Sledgehammer on Peter Gabriel's classic 'So.' The entire bass track seems to include the bass plus the bass dropped an entire octave.And the octave down bass line is mixed right up there with the original bass. Nothing subtle about it.

You can even use a pitch shifter to add two-,three-, or four-part harmony if you are so inclined. But get out your arranging book, because the

This results is a harmony or countermelody line with all the harmony and dissonance you desire. It's built on a single vocal track, and relies almost entirely on good sounding pitch shifting.

Go beyond harmonies. Use pitch shifting to turn a single note into an entire chord. String patches can sometimes be made to sound more orchestral with the judicious addition of some perfect octave and perfect fifth pitch shifting (above and/or below) to the patch.

NUTS and BOLTS

And don't stop with simple intervals.Chords loaded with tensions are okay too, used well. Yes put it front and center in 'Owner of a Lonely Heart' on the album 90125. Single-note guitar lines are transformed into something more magic and less guitar-like using pitch shifters to create the other notes.

A final obvious pitch shifting effect worth mentioning is the stop tape effect. As analog tape risks extinction, this effect may soon be lost on the next generation of recording musicians.

When an analog tape is stopped, it doesn't stop instantly; it takes an instant to decelerate. Large reels of tape, like two inch 24 track, are pretty darn heavy. It takes time to stop these large reels from spinning. If you monitor the tape while it tries to stop (and many fancy machines resist this, automatically muting to avoid the distraction this causes during a session), you hear the tape slow to a stop. Schlump. The pitch dives down as the tape stops.

This is sometimes a musical effect. And it's not just for analog tape, as Garbage demonstrates via a Pro Tools effect between the bridge and the third Chorus of 'I Think I'm Paranoid' on their second album.

Surgical effects

Pitch shifting is also used to zoom in and fix a problematic note. We've all been there.

In the old days of multitrack production (about a year ago), we used to sample the bad note. Then we tuned it up using a pitch shifter. It was raised or lowered to taste. Finally, the sampled and pitch shifted note was re-recorded back onto the multitrack. With the problematic note shifted to pitch perfection, no one was the wiser.

That was then. Now, clever software taking advantage of powerful hardware permits pitch shifting to be done automatically (Antares AutoTune [hardware and software versions], Wave Mechanics Pitch Doctor, TC Electronic Intonator, etc.). That is, the effects device can monitor the pitch of a vocal,violin, or didgeridoo. When it detects a sharp of flat note, it shifts the pitch automatically by the amount necessary to restore tuning. Wow. And it really works.

But please be careful with these devices. First,don't over polish your product.Pitch shifting everything into perfect tune is rarely desirable. Vibrato is an obvious example of the de-tuning of an instrument on purpose.

And if Bob Dylan had been pitch shifted into perfect pitch, where would folk music be now? There is a lot to be said for a musical amount of 'out-of-tuneness.' Remove all the warts, and you risk removing a lot of emotion from the performance.

> If Bob Dylan had been pitch-shifted into perfect pitch, where would folk music be now?

Second, don't expect to create an opera singer out of a lounge crooner. There is no replacement for actual musical ability. If the bass player can't play a fretless, give her one with those pitch-certain things called frets. If the violin player can't control his intonation,hire one who can.

Don't expect to rescue poor musicianship with automatic pitch correction. People want to hear your music, not your effects rack.

Out of time

This month represents our fourth month of discussion on delay. Are we done yet? Naturally, no. We continue our tour of the delay in a future Nuts & Bolts installment when we take a detailed look at reverb.

Alex Case strapped his pitch shifter to his gear shifter and drives by chanting.
Request Nuts & Bolts topics via case@recordingmag.com.