

Software Pitch-Shifter Comparison

Max Harrold

UIUC Physics 406

Spring 2016

Introduction



Electronic sound manipulation is everywhere. Anything from Alvin and the Chipmunks to modern hip-hop are good examples of things that would scarcely exist without it. These processes can include pitch-shifting, time stretching, pitch correction, and almost anything else you can do with a little technology to make something sound the way you want. As someone who is fascinated with music but doesn't really play an instrument, it was also the perfect topic for me to study this semester. I decided to focus on pitch-shifting, not only because there are many programs specializing in it, but also because it would have the most straightforward way of "fact-checking" these programs' claims.

Method

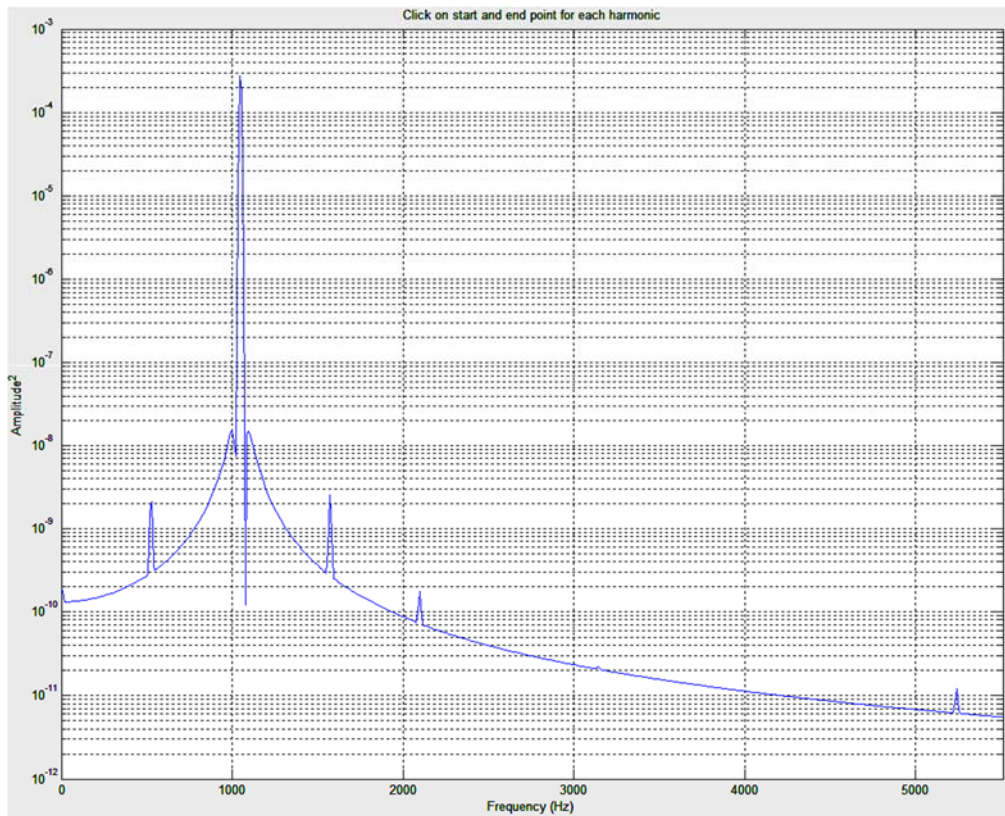

The digital audio workstation I decided to use was Logic Express, as I already was very familiar with it and I wanted to learn things that I could put to use in the future. Logic Express is the lighter version of Logic Pro, Apple's answer to professional recording software like ProTools. It's sort of the happy medium between Garageband and professional-level software. It comes with two built-in pitch-shifters, named "Pitch Shifter II" and "Vocal Transformer". In addition, I did some research and found a software plugin called "Melodyne", which was very expensive and touted as the best program to use.

I decided that the simplest, most replicable way to conduct these comparisons would be to pick a few software instruments, the sources of the sound, and play the same note in every test. With a very simple software synth ready, I played a ten-second long C4 as a sine wave, a square wave, and a somewhat cheesy preset called "Big Trance Now". I then played a C3 and a C5 with each instrument. These would be my control group: the sounds each pitch-shifter would be trying to emulate. With those ready, I then loaded the three C4s into each pitch-shifter, and moved them all up an octave, then down an octave. With all of these audio files ready, it was time to look at the frequency charts and see what had happened.



Pure Sine Waves, C3 (top) and C5 (bottom)





Melodyne Sine Waves, C3 (top) and C5 (bottom)





Pitchshifter II Sine Waves, C3 (top) and C5 (bottom)





Vocal Transformer Sine Waves, C3 (top) and C5 (bottom)





Pure Square Waves, C3 (top) and C5 (bottom)




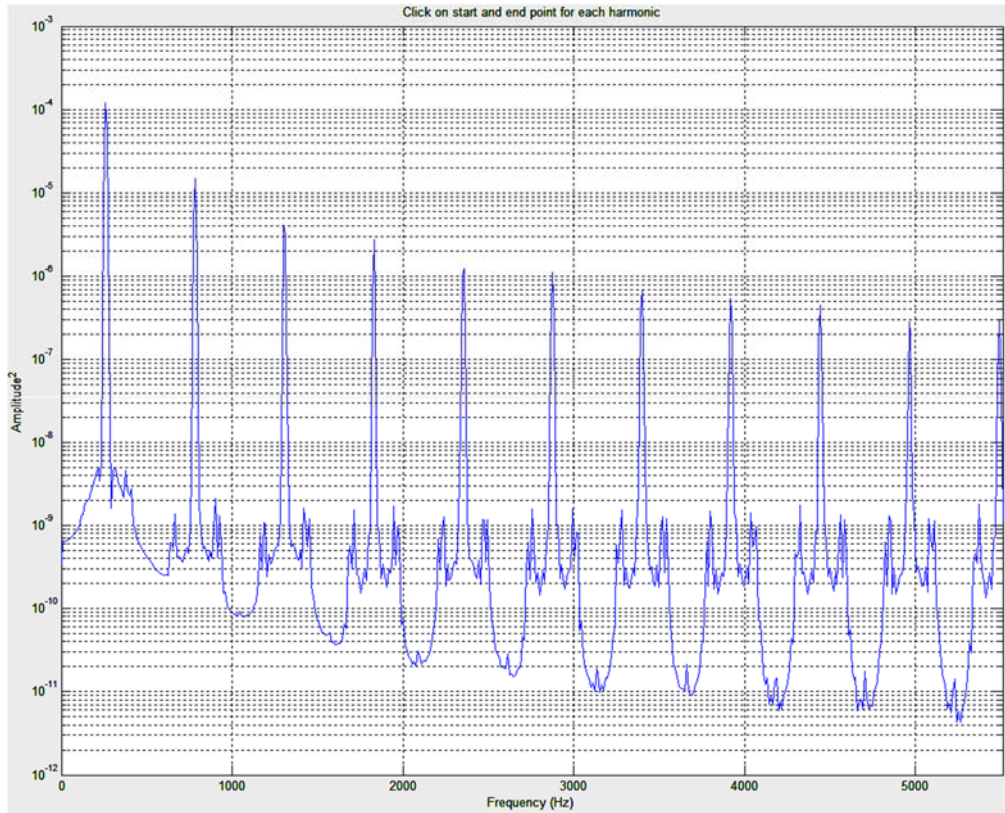
Melodyne Square Waves, C3 (top) and C5 (bottom)





Pitch Shifter II Square Waves, C3 (top) and C5 (bottom)





Vocal Transformer Square Waves, C3 (top) and C5 (bottom)





Pure “Big Trance Now”, C3 (top) and C5 (bottom)





Melodyne "Big Trance Now", C3 (top) and C5 (bottom)



Pitch Shifter II "Big Trance Now", C3 (top) and C5 (bottom)



Vocal Transformer “Big Trance Now”, C3 (top) and C5 (bottom)



Data

Even without getting very technical, it is clear from the frequency spectra which tests each program struggled with. The three instruments were each more sonically complex than the one before, and the data gets increasingly varied in response. The simplest sound, a computer-generated sine wave (with no distortion from recording equipment or room acoustics) was the easiest for all three programs to handle. That makes sense, because if I was to build my own pitch-shifter, sine waves are where I'd start.

The only graph that is not so straightforward to interpret is Pitch Shifter II's version of the Big Trance Now C5. The peaks and valleys are all where they belong, with relatively consistent amplitudes, but there is a surprisingly clean and clear interference pattern in the valleys that only appears in that test.

Shortcomings

Unfortunately, this test did not account for why each program performs the way it does. I would still like to learn how each program interprets and manipulates the audio it's given, although there is a good chance that that is more in the world of computer programming than the physics of music. Not only that, but I imagine in some cases, Melodyne especially, that it is proprietary information that the average user would not have access to.

From a purely economic standpoint, Melodyne can probably get away with being second best in pitch-shifting because of all the other effects it features. It is the only one of the three that does anything more than pitch-shifting, and it really does do a lot more.

Another consideration is the manipulation of non-computer-generated audio. Would Melodyne do best with live guitar? Maybe. Would Vocal Transformer do best with human vocals? Probably not. It would be very interesting to see, however.

Conclusions

Interestingly enough, across all tests, Melodyne (the professional, expensive one) did not do the best. Logic Express' own, poorly named Pitch Shifter II was the truest to the source material every time. Melodyne was a close second, and Vocal Transformer did very badly. The cool thing is, you can hear these differences too. Especially with Vocal Transformer, the pitch-shifting artifacts are audible and distracting. And how things sound is, of course, the most important thing.