


Third Octave Filtering



One of the familiar sounds of the early days (50s & 60s) of electronic music came from 1/3rd octave filters. This type of device consisted of 20 to 30 sliders on a heavy box. Each slider controlled a narrow band of frequency, spaced at 1/3rd of an octave. Half octave and octave spacings were available (and cheaper), but the 1/3rd octave design was the Cadillac.

Here's how to get that sound in MSP.

The classic version:

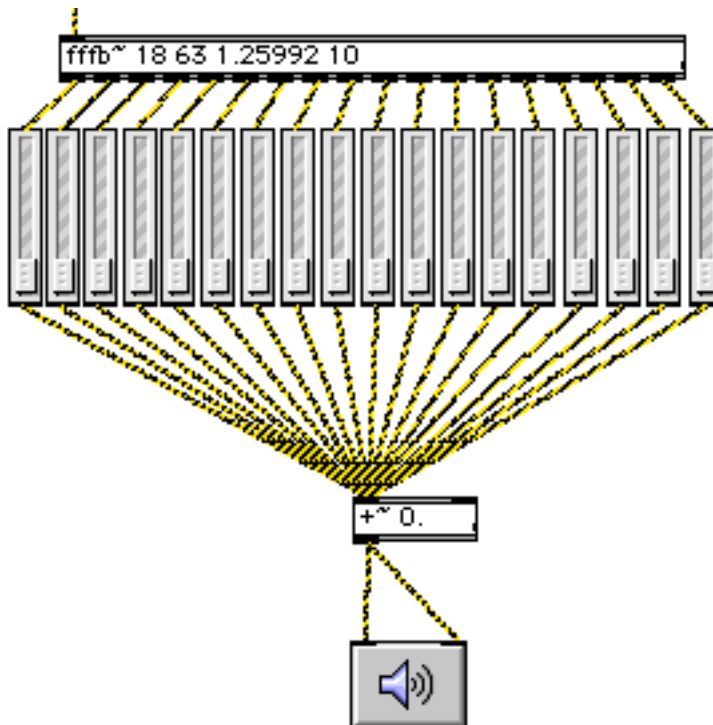


Figure 1.

Figure 1 illustrates the use of `fffb~` to work at 1/3 octave spacings. This one has 18 bands, starting at 63Hz with a spacing ratio of the cube root of two. The Q^1 of 10 gives the filter a sound similar to the models I remember. These are not equalizers- EQs usually have a gentler slope in each band, and are designed to give a nice flat response when all of the sliders are at the mid-point. 1/3rd octave filters are analysis tools for isolating frequency components. They had a Q of about 4 and a cutoff slope of -48dB per octave or so. `fffb~` has a slope of $12\text{dB}/8\text{va}$, but the high Q gives a good cutoff effect, with a slightly narrow passband. This is typical of some low end filters made primarily for electronic music. They always had a distinct chordal sound with the sliders set evenly. You can get more pristine sound by doubling up the `fffb~`, using two or even three sections to increase the ultimate cutoff. This lets you reduce Q for a flatter passband, but for most electronic music effects, a single section with high Q is fine.

Note that the individual outputs are summed. This is the way these machines came from the factory, but many studios modified them to get individual control of the bands. In my days at the University of Iowa electronic music studios, I built a keyboard controlled version (after a design by Lowell Cross). The equivalent is shown in figure 2.

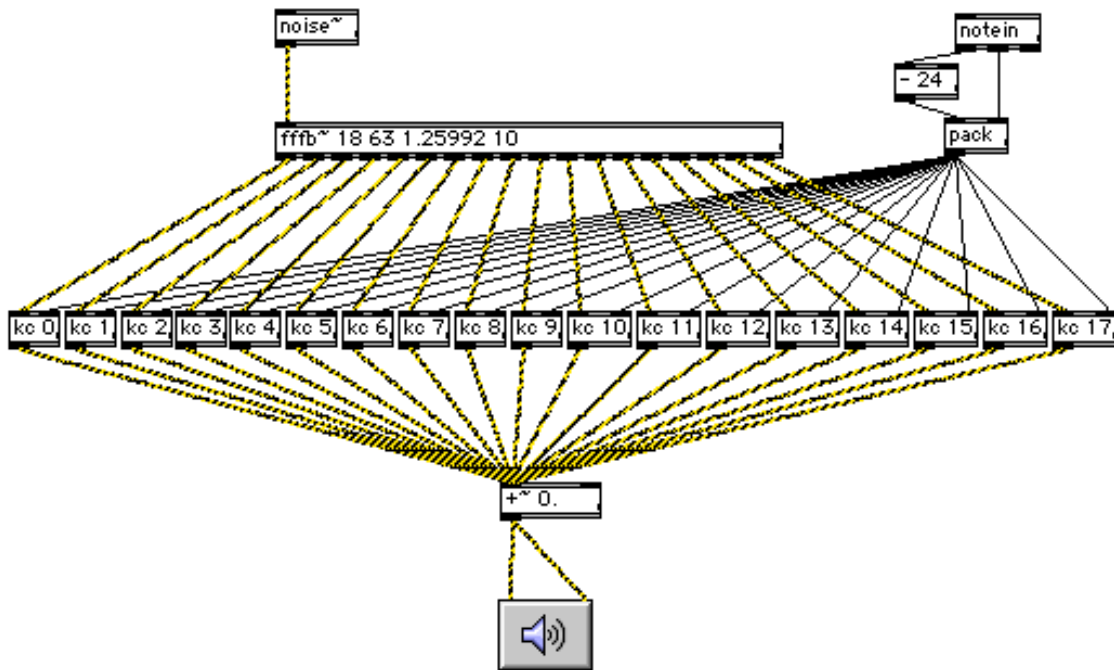


Figure 2.

¹ To review, Q is the ratio of frequency over bandwidth. Bandwidth is defined by the frequencies where the signal is reduced by 3 dB.

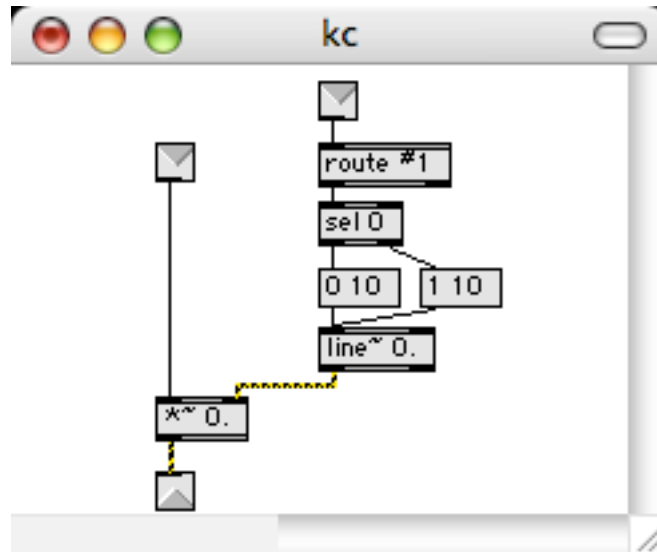


Figure 3.

Figure 3. shows the contents of the control subpatch. Note the use of #1 to allow a different routing for each band.

Of course this design does not take advantage of the programmable features of `fffb~` itself. In figure 4, a multislider generates messages to control internal gain of the `fffb~`




Figure 4.

The results are much the same as figure 1, except the faders are now linear instead of logarithmic. This is fixed with a `vexpr` (and some redefinition of the multislider, which now has a range of -60 to 10). I've also added a control to adjust the Q .

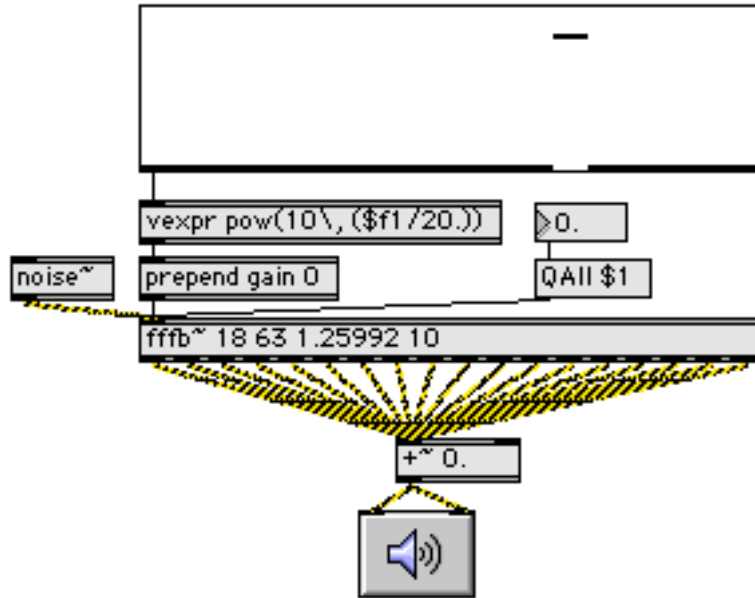


Figure 5.

As I mentioned, doubling the fffb~ gives a purer effect. Compare the output of figure 5 with that of figure 6. The difference becomes more obvious as you increase the Q.

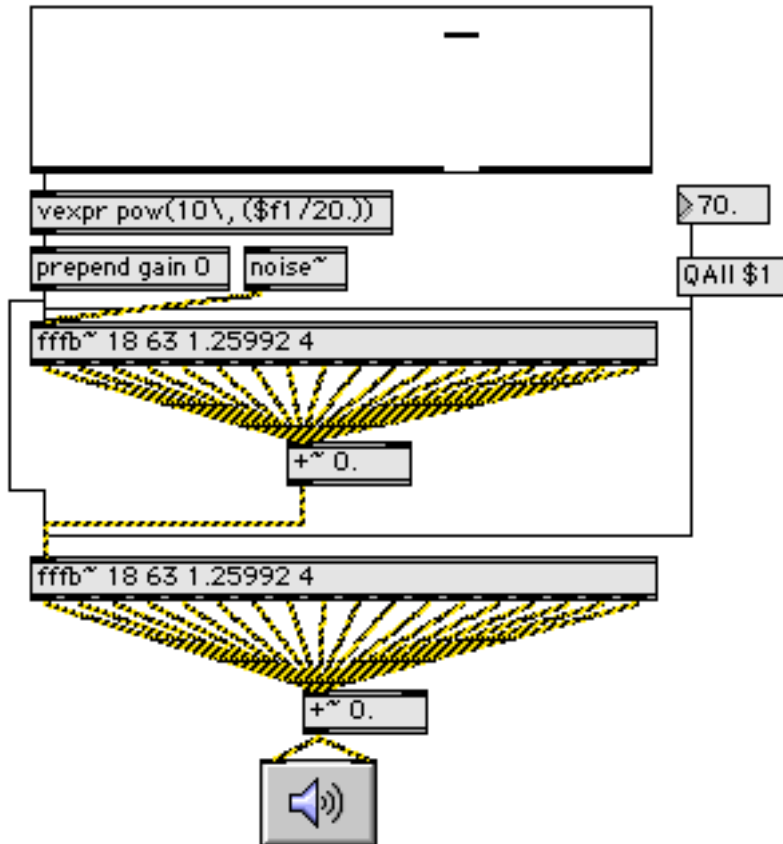



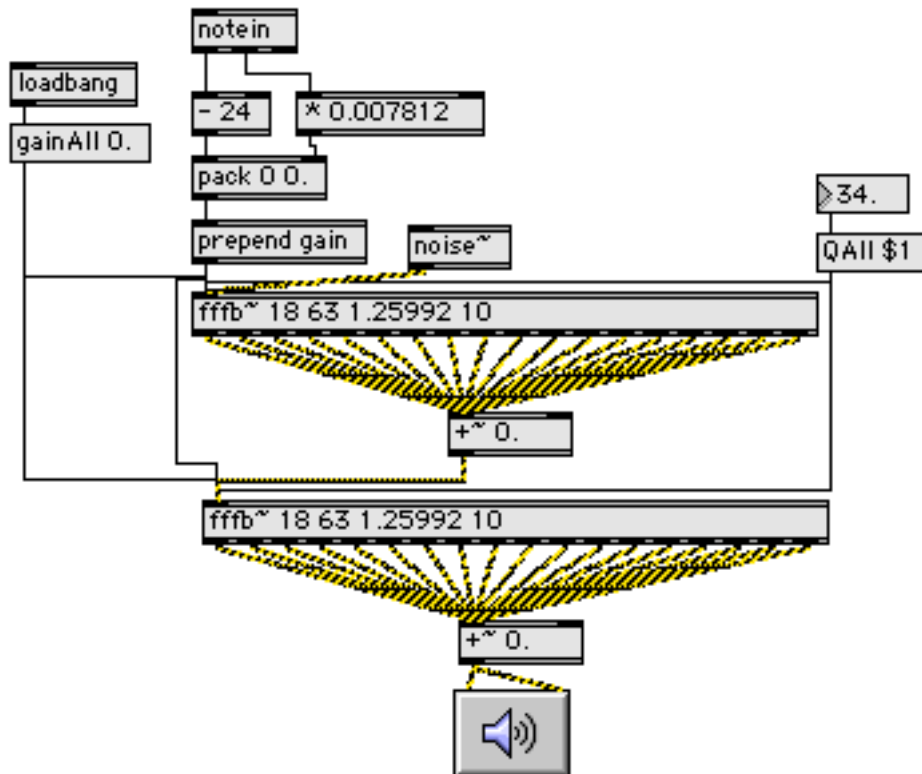
Figure 6.



With higher Qs the double circuit gives a very clear tone. The spectrograms in figure 7 illustrate this.

Playing tones

A keyboard version is shown in figure 8. It's interesting to experiment with different frequency mappings, major chords in just temperament, for instance. All that is necessary is to load a list into the filters as demonstrated in the help patch².



² One of my favorite sets is used in Buchla synthesizers. The frequencies are 100, 150, 250, 350, 500, 630, 800, 1k, 1.3k, 1.6k, 2k, 2.6k, 3.5k, 5k, 10k. Another interesting set is the Bark scale from psychoacoustics: 50 150 250 350 455 570 700 925 1175 1375 1600 1860 2160 2510 2925 3425 4050 4850 5850 7050 8600 10750 13500.

Figure 8.

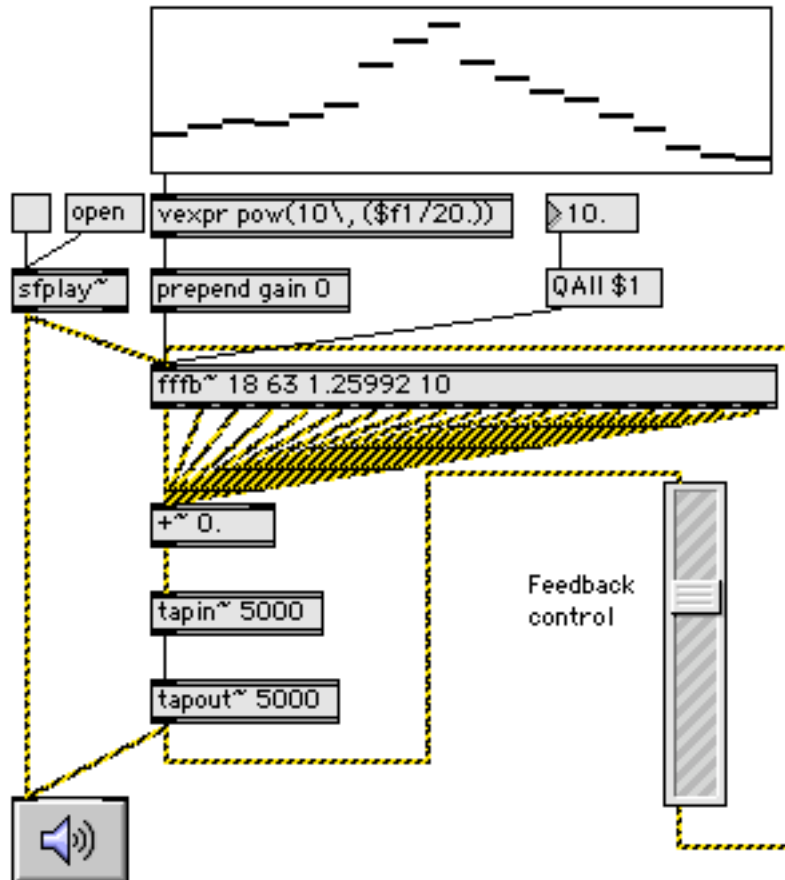
The Old Feedback Trick

Figure 9.

One especially lovely technique of the old days was the "screen". This was tape delayed feedback³ through a 1/3 octave filter. Each pass through the filter would accentuate the filter's effect, ultimately resulting in something like a diminished chord. The patch in figure 9 will do this. You only need to play 5 or 10 seconds of sound- the feedback will take over from there. The longer you go, the purer the tones.

Vocoders

The final chapter in this look at multi-filter applications is vocoding. The name comes from "voice encoding" a technique the phone company pioneered as a data compression scheme. Vocoders use two multi-band filters; one that analyzes the incoming signal, and one that uses the analysis to impose the spectrum on a second signal. Figure 10 shows how this works:

³ With two tape decks, the tape threaded across the table from one to the other. Alvin Lucier's "I AM Sitting In A Room" does this, using room acoustics for the filter.

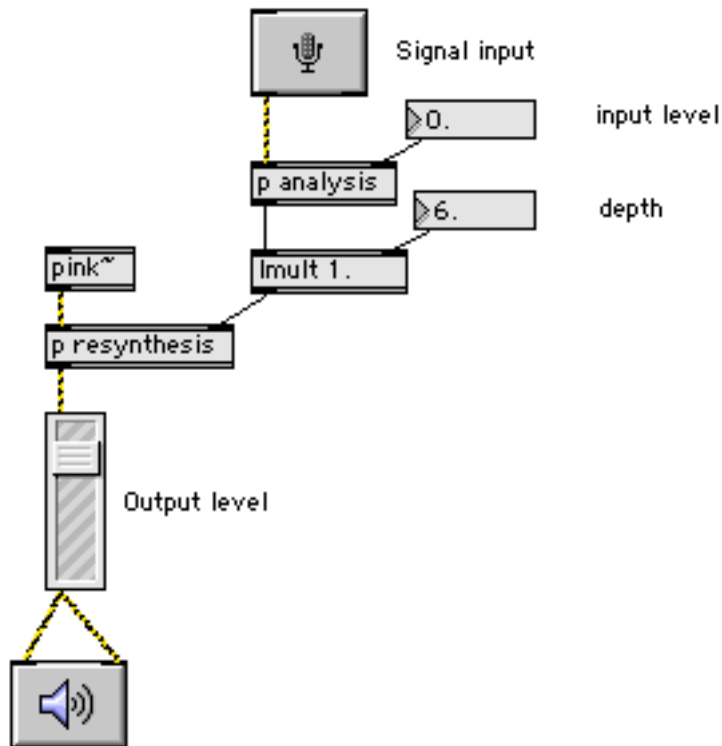


Figure 10. Filter based vocoder

The voice comes in at the ADC~ object and straight to the analysis subpatch:

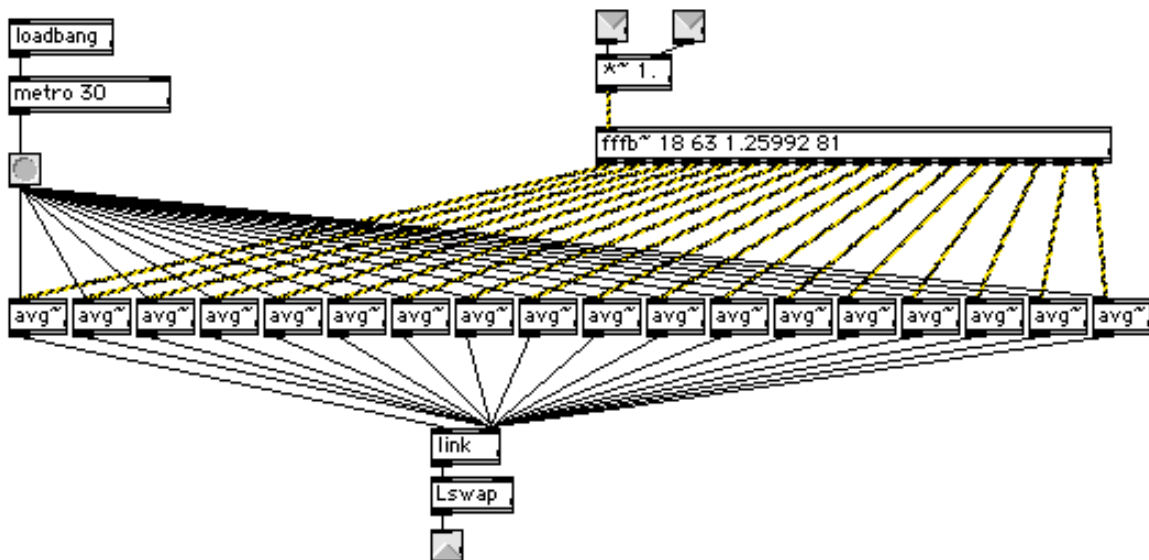


Figure 11. Analysis

The analysis is done by the avg~ objects, one per filter band. These numbers are grouped into a list, which will be backwards (right to left action) but flipped around by Lswap.

The list is used to control the resynthesis as shown in figure 12. As you can see, the analysis section replaces the multi-slider used in early examples.

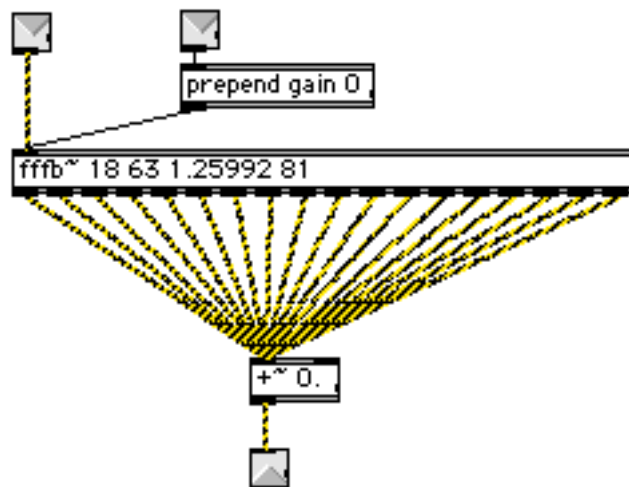


Figure 12. Resynthesis

The resynthesis filter should match the analysis filter if you are trying for clear speech, but other frequency mappings can be entertaining. The Q of the analysis filter can be adjusted for various types of input material. A two stage resynthesis filter will let more signal through at low Q and give clearer tones with high Q settings.

When voice is the input, you can almost understand the speech, but all you are getting is the vowel sounds.⁴ Intelligibility can be improved by bypassing the vocoder through a high pass filter, mixing some sibilants and consonants into the output.

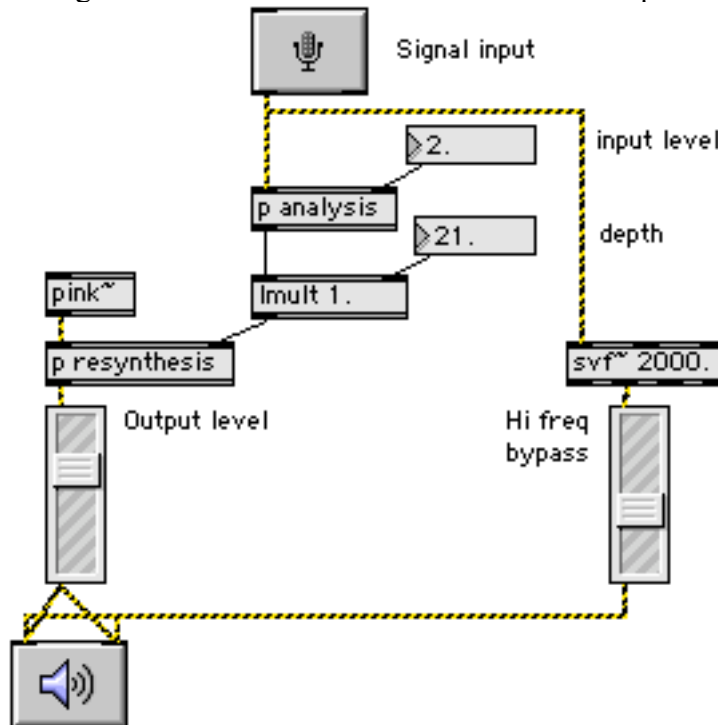


Figure 13. High frequency bypass.

⁴ Probably why vocoding never caught on in the communications business.

A pitched instrument can be made by substituting any number of synthesis patches for the pink noise input. Figure 14 shows an instrument that produces results similar to the Bode vocoder Wendy Carlos used in the movie "Clockwork Orange".

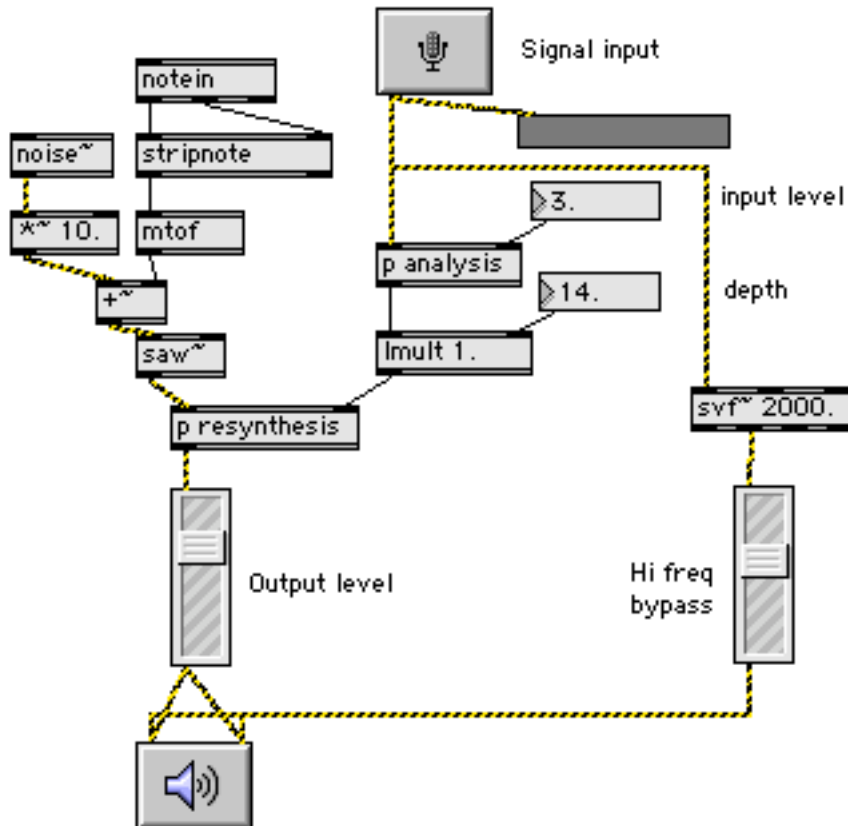


Figure 14. Pitched vocoder

The foundation for the resynthesis side is a sawtooth wave modulated by noise. This produces a spectrum that is pitched but with components spread across the bands of the filter. (Without this, the output is uneven when the saw plays pitches in the cracks between bands.) The resynthesis filter has a lower Q than before, around 30.

Multiband filters are used in a number of other places—the newest being multiband compressors (or "optimizers") used to process nearly all commercial music. Of course the ultimate multi band device is the fast fourier transform (see `fft~`), with divides the audio spectrum into up to 2000 slices.