Frequency Tweakers 1

A variety of filters and frequency based effects

Cupwise Nebula Programs

General Information

The main theme of this series is that all effects in it will somehow act upon the frequency spectrum- as equalizers, filters or filter based effects. Some unusual hardware (stuff you wouldn't find in studios) will always be involved in the sampling process, either to perform the freq spectrum manipulation, or to add color (tone) to the final effect. Other key elements of the series will include experimental ideas/designs, and final products that really push what's possible with Nebula by using more samples, or using the samples in new ways, to create more in-depth effects. These effects use new custom sampling templates to do things that just haven't been done before. All will include at least some dynamic samples, which isn't typical with these kinds of programs. Most if not all Nebula EQ and filter effects to date have been 'static'. Doing them with dynamics requires huge amounts of samples per program, more CPU/RAM to use them, and much more time is spent making them! In some cases you can expect longer than usual loading times, unless you have a very current system. As always there are lower kernel options to use which tax system resources less. Instead of talking in general terms about these ideas any further, let's look at these first two effects to get an idea of what I'm talking about!

Installation:
Easy! You just place the .n2p files in your Nebula programs folder, and the .n2v files in the vectors folder. You will find the 'Tube Radio LP' programs in the 'FLT' category in Nebula, then in the 'FT' sub-category for the 96khz versions, and the 'FT*' sub-category for 44.1khz. The 'Tube Radio EQ' programs are in the 'EQU' category, then in 'FT' for 96khz, and 'FT*' for 44.1khz.

Tube Radio LP:

This is the most complex synthesizer style low pass filter done with Nebula to date!

- Smooth control of the cutoff frequency, from below 100hz to around 17khz.
- Fully adjustable resonance control can go from a smooth roll-off, to a very sharp and prominent resonant peak.
- 4 and 10 kernel options. 10 kernel version can get nasty.
- To push the envelope even further, it was sampled with a (very) small amount of dynamic behavior!

Dynamics in a Nebula effect like this is unheard of. The reason for it here is very specific. I wanted a low pass filter with tube distortion possible, but with some ability to shape/control it. At lower input levels you might not notice the harmonics/distortion, but boost the input and it will get more and more grungy, especially with resonance at 100.

Here's how this thing was sampled- first the signal went out D/A, and into an A-Station synthesizer which provided the low pass filter (12db/oct). From there it went to an old Bogen PA amplifier (model CT-60) to drive the next stage (and give a little flavor). Then it went to an EDM FM radio transmitter, and the signal was picked up by the Bell FM tuner (so radio waves were part of the sampled chain here). Now, the key to this is that the resonant peak was squashed down drastically, for the higher dynamic portion of the effect. The Bogen amp was used to drive the signal into the transmitter/radio at JUST the right level, so that only the resonant peak would be squashed and only when high resonance was used with the filter. The distortion generated came from both the transmitter itself and the tube radio.

The end result is an effect that can go from having a very fluffy, smooth sound with lower input levels, to having a growling sound you won't hear anywhere else. A hot signal + lots of resonance gives you higher distortion levels, with a strong focus on the resonant peak. So if you have the cutoff at 3khz, you'll get strong harmonics at multiples of that (6,9,12,15khz,etc). The harmonics of that resonant peak occur above the cutoff point, where the input audio has been strongly attenuated. So anything above the cutoff point is cut as usual, but is replaced by distortion from the resonant peak. Of course, this is only if you want it to work that way. The 'trim' control can be used to dial in as much, or little, of that behavior as you want.

There are a few different variations offered. There are 2 different kernel options- 4 and 10. The 4k version also has its h2-h4 harmonic samples trimmed to a slightly shorter length which also helps with CPU/RAM (but gives a little less

detail as with the lower kernel count). Then there is a special 'HQ' program, which is the 10k version but with the fundamental and even harmonics set to use 'timed' mode instead of 'freqd'. This sounds a little better, but takes a LOT more cpu, and this is an effect that already uses a fair amount so be warned. The program rate is also set to be faster here, which may make filter sweeps a little smoother in some cases (and also costs more CPU). The HQ version is intended for use on render, to get the absolute best quality. If automation is recorded for the cutoff control with a non-HQ program in any host, that automation should still be the same when the program is switched to an HQ version, so you could use the standard versions for working and then switch to HQ mode before render. Any adjustments made to the other parameters would have to be carried over after switching to the HQ program (unless they were automated also). See tips section for more on this.

Controls:

Cutoff- What you would expect it to be in an LP filter! If you max this out, you will still have a low pass filter at around 17khz. Because of the way the FM modulation and demodulation process works with radios, everything needs to be cut above ~17khz before a signal can be transmitted. So at max cutoff you don't hear the synth filter any more, but the filter from the transmitter.

Resonance- Lowest setting gives no resonance, highest settings gives a very sharp peak of around 15db (depending on cutoff freq). High res settings also slightly bring down the overall level of the signal in the pass-band, as a kind of compensation for the resonant peak. Highest resonance settings trigger the most distortion (as explained above). At higher frequencies you will lose the resonant peak, because it was filtered out by the FM transmitter.

Attack & Release- The same controls you find in most preamp-style Nebula programs. They are here because there is some dynamic action in the effect, so these parameters give you a little more control over the final result. They have been set to default positions which will work well in any given situation, so you can leave them alone if you want. On the other hand, you might play with them to see if you can get different results in some cases. Faster attacks and slower release settings will tax CPU more, and may cause glitchy behavior.

Trim- The famous Nebula trim feature. It's here because this isn't just a filter, it's a DYNAMIC filter. To see the full range of this filter, use a 10k version and try sweeping the cutoff control up and down over some audio, with the resonance at 0. Then raise res to 50 and sweep again. Then again with res at 100. Now slowly raise the trim control. Be careful, if you go too far Nebula will spit out bad distortion that you don't want, but if you go slowly you can hear it starting to happen- at that point back off a little. Now sweep the cutoff again at different

resonance settings. To go one step further (and get extra nasty), raise the 'dist' control up a little and sweep again.

Dist- The classic 'drive' Nebula control. Cupwise programs usually rename it to 'dist' because it's a little more accurate of a description. This control is basically a mixer for the harmonics generated by Nebula. For the most natural sounding results, it's probably best to crank trim BEFORE increasing dist level, but if you want the grittiest that this filter has to offer, you will need to boost both.

Liquidity- This being a filter program, and filter programs being made to be swept around, liquidity might be useful here. More liquidity should smooth the sound of the filter sweep. Try different amounts of it and sweep the filter around to see if you like what it does. A little liquidity is used by default on loading the programs, so you can even try setting it to 0.

TubeRad Harmonic EQ:

A completely unique Nebula EQ/distortion tool. Nothing like this exists anywhere else. An adjustable 'radio band' provides harmonics to a frequency range selected by the user, and can be boosted/cut to provide some EQ.

- 3 separate modules- low, mid, and high. The 'radio band' for all 3 modules was processed twice (two passes) by a Bell tube radio tuner. The mid and high modules also included an old Philco tube amp taken from a phono console, in the 2nd pass.
- Each module allows you to position a 'radio band' somewhere in the frequency spectrum, and further define it with 'width' and 'gain' controls, just as you'd find in an EQ. The band will give some analog color to just that area, everything outside of the radio band is dry. This can be used to provide fuzzy distortion or subtle harmonic excitation to a specific part of the audio.
- Low/high modules are shelves with variable frequency settings. Mid module is more like a 'peak' band, and allows user to pick one of 5 center frequencies.
- Also dynamic, so different input levels trigger different distortion content and dynamic behavior, as sampled from tube radio/amp setups. Combine this with the ability to mix exactly how loud the harmonics will be, and you have a very tweakable precision distortion tool.

This one is actually very experimental in the sense that it never existed outside of Nebula, and can only exist inside it (it was tailor made for the platform). It's a hand-crafted EQ/effect. The main idea for it came when I noticed that some of the Cupwise 'Tube FM 1' effects can be used as subtle harmonic exciters/enhancers, but I wondered if it would be possible to have more control over that ability. I wanted a frequency specific distortion/harmonic exciter. I had an idea for 'stitching' together an effect where the user could position a single band at various frequencies, with control over the band's width, just like an EQ. The adjustable band would be made from signal that was processed by a tube radio (and would thus produce harmonics), whereas everything else would be totally 'dry' (producing no harmonics).

Almost a year after I first had the idea, it's now a reality. It's split into 3 parts ('modules'), one each for low/mid/high frequency ranges. High and low are shelves, while mid is a peak. The 'radio band' can be adjusted by around +-13db (it depends on which freq 'module' you are using, but the Nebula readout always goes to +-13db because it really can't be accurate since it even depends on the freq selection), giving this effect the added ability of functioning as an EQ. The width can be adjusted from very sharp (except with the low 'module', which doesn't get so sharp) to fairly wide. No EQ was sampled to make this effect, it was all designed 'by hand', through careful mixing of the radio processed signal (the tones NAT uses in sampling) with the dry. The slopes of the bands were made from custom amplitude envelopes applied to the NAT tones (in Reaper). So the EQ aspect of this effect was created digitally (digital gain adjustment of the radio band), but is handled by Nebula, so it sounds very good. Because the adjustable band was processed by tube radio, the final effect sounds very analog, while also having qualities similar to linear phase EQs.

The ability to actually adjust the radio band's gain relative to the dry band (and thus use the effect as an eq) came as an afterthought. The main point of this effect is to be able to position the radio band at different points on the freq spectrum, then input audio will generate subtle to not-so-subtle harmonics, but only for content that goes through the radio band. The traditional Nebula drive control (renamed 'dist' here, as with other Cupwise effects) can then be used to decide exactly how audible those harmonics will be. You get the harmonics even if you don't boost or cut with the band, so the effect should be thought of as a 'frequency specific distortion/harmonic exciter that can also equalize' or something along those lines.

On top of everything, this effect is DYNAMIC. In Cupwise tradition the envelope has been pushed here, making this almost surely the first and only EQ with 3 controls AND dynamics. The higher dynamic levels that were sampled resulted in the radio being over-driven, and a little compression came from that.

Each dynamic step of the 'dry' tones had to be level matched to the corresponding radio sample, to get a coherent effect (but the gain adjustment is the ONLY processing ever done to the dry tones). This means that not only can you decide exactly how audible you want the distortion to be by using the 'dist' control, but you can also get different sounding distortion/dynamic behavior by using the 'trim' control to drive the effect differently.

The Bell tuner was used to make this effect, mainly because of it's almost totally flat frequency response, but also because of it's phase response being nearly identical to that of a 'dry' signal, which is why this thing is possible. It was sampled without FM transmission- the signal was directly 'injected' into the radio's circuitry. There were actually 2 passes through because the first time didn't generate enough harmonic content. For the 2nd pass, the signal also went through an old Philco tube amplifier (which will be featured more in the 2nd part of this series), but only for the mid and high freq modules (for the low module there was only the 2 passes of the Bell).

QUICK SUMMARY- It's like an almost-linear-phase EQ with some analog qualities, or like a frequency specific harmonic exciter with the ability to do EQ (which is usually close to linear phase for the mid and high modules). Use the 'low' program/module if you want to add some harmonics only to the lower frequencies of your signal (and not the rest). Use mid or high to add them to those parts of the signal. You get the harmonics regardless of whether you boost or cut, so the 'freq' and 'width' controls still work to define the area that they will be generated by, even with the band's gain at '0'.

Controls:

Freq- Selects the cutoff frequency for the shelve bands in the low and high programs, and the center frequency for the peak band in the mid program. If using the low program, harmonics will be generated in the frequency range below where you set this control. For the high program it's the opposite. For mid, harmonics will be generated at and around the frequency you select. This control is fully variable for the low and high modules, but for mid it switches between 5 frequencies.

Width- Adjusts the steepness of the slope(s) for the radio band. For low and high you are adjusting the roll off of the shelf, for mid you are adjusting two slopes, one on each side of the selected frequency. 0% give you the steepest slope, 100% gives widest. This control doesn't always work as you would expect it to, but there is no loss in quality because of that. You just get unusual shapes sometimes (see notes below for explanation).

Gain- Adjust the gain of the radio band. The results are usually not symmetrical

above and below 0db (the shape can be different).
Freq, Width, and Gain controls are just like you'd find in an equalizer. Just remember that the band they are used to define is special here, though. Anything that goes through that band will generate tube radio harmonics. The accuracy of these controls is not perfect, so you often won't get exactly what you expect due to your settings. You have to use your ears. It might help to raise 'dist' level so that the generated distortion is clearly audible while adjusting where you want it to be and how wide you want the band to be, then you can lower 'dist' back down to where you actually want it.
Trim- You can use this to adjust the level going into the effect, with compensation for the output. It will change the harmonic content, because this is a dynamic effect.
Dist- Classic Nebula 'drive' control, renamed. Use this to directly control the mix level of the generated harmonic content. Since you should be using the effect as an 'exciter', this control should come in handy here.

NOTE- I would usually recommend using Vst Analyser to see what kind EQ shapes you will get with this, but for this particular effect it sometimes shows false results (I think because of the complexity of the effect- the dynamics throw off the analysis). The analyzer in NAT seems more accurate. At the end of the day, an analyzer will only get you so far, and you always have to use your ears. If you do check these out in an analyzer, you may notice that you can tell in some cases where interpolation is being used to transition from one sample to the next. It's most noticeable with the width control during a boost with the high module. This isn't such a bad thing here, because this EQ wasn't made by sampling an analog EQ, so having super precise transitions isn't necessary. The EQ side of this effect is a purely digital effect so larger amounts of interpolation won't hurt the quality or accuracy, because it isn't 'modeling' anything. It would be nice to have smoother transitions with the width shape, but because of the dynamics, it would just require an insane amount of samples in the end and it isn't worth it. So, you can get some unusual shapes between the extreme width settings, but I felt it was better to keep the ability to interpolate between them instead of using switch mode. Besides, in some cases width shapes actually transition relatively well.

General Usage Tips/Ideas:

- The Radio LP effect can take lots of CPU. The HQ version will give the best possible sound quality and the smoothest filter sweeps, but it takes more cpu. Switching a program in Nebula shouldn't erase or undo any automation that has been recorded with Nebula controls, so if you use the non-HQ version to record the automation, then switch over to the HQ mode just before render, it should give you the same filter sweeps you recorded but with the better quality. The only problem is that any adjustments you made to the res, trim, and dist controls WILL be lost when you switch to HQ, unless you automated those also. It is up to you to figure out how to deal with that, and it will depend on your host. It shouldn't be too difficult in any case. You could use the host's automation system to place those controls where you want them, without having any actual recorded movement. That should keep them in place when you switch to HQ. On the other hand, it may be quicker to just keep note of where the controls are and re-position them after switching.
- The Radio LP has 4k, 10k, and HQ options. Both have HQ and 'standard' versions. You can use the 4k version for the absolute lowest CPU/RAM usage. If you want distortion you should always use the 10k option, using standard to work, then HQ to render. If you don't care so much about distortion but just want the LP effect, go with the 4k option and lower input levels. There isn't an HQ preset for the 4k version, but you could just manually set the kernels to use timed mode yourself on the kern page. I would recommend only doing this for the fundamental and either the even or odd harmonics, but not both, because setting both even and odd to timed can cause artifacts. You don't need to worry about adjusting the 4k version to have the quicker program rate (like the 10k HQ has compared to the standard 10k), because the 4k preset already uses the quicker rate.
- Try this with any of the Radio EQ modules: load it onto a track with material that has a fairly complex signal, with a wide range of frequency content (for example, a mix). Use the gain control to cut several dbs, out of a particular frequency area. Now, 'compensate' for what you just cut by adjusting the 'dist' control, boost it by however much you cut by, or more. You are sucking out some of the fundamental and replacing it with harmonics.
- Don't forget the importance of the trim control to these effects. Sampling these effects with dynamics makes them much more difficult and time consuming to make. So put those dynamics to good use! Trim allows you to quickly find the sweet spot.

Thanks:

Giancarlo, Enrique, and everyone else at Acustica!
www.acustica-audio.com

Zabukowski- maker of Nebulaman. I used this program a little, as a tool, during the development (some) of these effects. Its function is to apply one or more Nebula programs to multiple audio files. It's a batch processor for Nebula. It did the job MUCH faster than a batch script with my favorite .wav editor! You can use it to process each track in a mix with a favorite Nebula effect or chain, to add a little analog flavor to each element before mixing. Check it out!
http://www.zabukowski.com/software/

Thanks to my customers/supporters, and everyone who has helped with and participated in the development of the Nebula technology.