NastyVCS

MANUAL

revision 1.0

Content

1 Introduction

1.1. License

1.2. Installation

Requirements:

- Win32 compatible system with SSE2 (or higher) instruction set support
- Tested and known to work in many VST compatible hosts

Put the DLL file contained in this archive in the VST plug-in folder of your host.

1.1. Overarching topics

Warning: Lower your listening volume while operating the plug-in to avoid hearing damage or damage of speakers or any other equipment.

Usage tips:

- Use the power switch on the right side for handy A/B comparisons
- In the toolbar in the bottom of the interface the IN switches must be active to run each section - use them for A/B comparisons as well
- Use <ctrl> + mouse left click on a knob or switch to restore default position
- Use <shift> + mouse left click on a knob to fine adjust values
- Use this plug-in as an insert effect in any mono or stereo channel of your VST host

1.1. Credits

The outstanding visual concept is done by Patrick Barca, www.suxesiv.ch. Thanks again, mate!

Many thanks to Philipp @ Torpedotrickser for so much time and support on setting up the compressor timings right.

Many thanks to all the beta testers.

2 Jump Start

2.1. Overview

NastyVCS – Virtual Console Strip.

Inspired by the smooth dynamic and tone shaping capabilities of some high-end mixing consoles and channel strips, this plug-in implements the most distinctive and much appreciated sonic effects generated by these devices:

- filtering and equalizing
- preamp style saturation and phase adjustments
- opto-electric compression

Functions at a glance

- performs gentle audio dynamic treatments
- masters difficult to handle audio material in a musical fashion
- shapes frequency and phase response
- adds extra harmonics and saturation effects
- controls outgoing audio peaks

Plug-in specification

- Win32 / VST compatible
- state-of-the-art digital signal processing
- performance-critical parts are written in assembler
- completely SSE optimized

Getting the most out of it

Please read the following chapters to get the most out of this device. Learn how to efficiently set and combine each section:

- obtain some effective tips on getting the most out of the presets
- understand the basic workflow of this device
- learn how gain-staging is handled
- take advantage of the sidechain and routing options
- see how dynamic treatments can be split across compressor, limiter and saturator
- understand how EQ and filtering is implemented

2.1. Quick reference

From left to right according the graphical interface :

#	GUI label	Label in host automation	Description
1	IN	Preamp In	Turns the preamp stage on
2	SAT	Preamp Sat Level	Sets the amount of preamp saturation
3	IN	Preamp Input Level	Sets the audio input level (post saturation)
4	IN	Filters In	Turns the two filters on
5	SC	Filters to SC	Routes both filters into the sidechain path of the compressor

6	HP	Filters HP Freq	Sets the HP filter frequency
7	HiQ	Filters HP HiQ	Switches into high Q mode (24dB per octave instead of 12dB per octave)
8	LP	Filters LP Freq	Sets the LP filter frequency
9	HiQ	Filters LP HiQ	Switches into high Q mode (18dB per octave instead of 12dB per octave)
10	EXT SC	External SC	Takes the compressors sidechain signal from audio input 3 and 4 (instead of the main channel 1 and 2 in's)
11	IN	EQ In	Activates the EQ section
12	POST	EQ Post Comp	Routes the EQ section behind the compressor
13	BOOST	LF EQ Boost	Low frequency boost
14	CUT	LF EQ Cut	Low frequency cut
15	FREQ	LF Freq	Low frequency center
16	HiQ	LF HiQ	Switches into high Q mode (results in steeper EQ curves)
17	GAIN	LMF Gain	Low mid frequency gain amount
18	FREQ	LMF Freq	Low mid frequency center
19	HiQ	LMF HiQ	Switches into high Q mode (results in steeper EQ curves)
2ß	GAIN	HMF Gain	High mid frequency gain amount
21	FREQ	HMF Freq	High mid frequency center
22	HiQ	HMF HiQ	Switches into high Q mode (results in steeper EQ curves)
23	AiR	AiR Gain	Turns in the high shelf
24	8k	-	Sets high shelf to around 8k
25	12k	-	Sets high shelf to around 12k
26	17k	-	Sets high shelf to around 17k
27	COMP	-	Sets metering to display compression gain reduction
28	LIM	-	Sets metering to display limiter gain reduction
29	OUT	-	Sets metering to display output volume
30	P-IN	Phase In	Turns the phase tool on
31	INV	Phase Invert	Switches the polarity of the audio signal
32	90/180	Phase 90/180	Switches between 90 and 180 degree phase shift
33	PHASE	Phase Shift	Adjust the phase shift from higher to lower frequencies
34	IN	Comp In	Turns the compressor on
35	+12dB	Comp Input +12dB	Boosts the compressor incoming signal 12dB
36	COMP	Comp Level	Sets the compression level (drive)
37	GRIND-PRESS-S-LACK	Comp Attack	Sets the overall attack time characteristics
38	SQEEZ-THRUST-RELAX-SOFT-	Comp Release	Sets the overall release time characteristics

	L.A.		
39	MAKEUP	Comp Makeup	Gain adjustment after compression
41	IN	Limiter In	Switches the limiter on
42	FAST	Limiter Fast Release	Selects a faster release time characteristic
43	LIM	Limiter Level	Turns the limiter transfer curve from 1:1 to ∞:1
44	OUT	Output level	Output level
45	-	Power	Turns the whole plug-in on/off

2.2. Basic operation and advice

Use this plug-in as an insert effect in any stereo or mono channel of your VST host. It can be operated both as a mono or stereo plug-in. If your host supports sidechain routing the external sidechain feature can be used as well.

Make sure that the power switch on the right side is in on position now (lightning). Make also sure that in the toolbar in the bottom of the interface each sections IN switch is active (lightning) if this section is needed.

For example turn on the IN switch for the compressor (straight below the COMP knob) and dial in some compression with COMP. Compensate for volume loss with the MAKEUP dial. The attack and release time switches can be dragged up and down to select a specific program.

Use the IN switches for A/B testing.

While the IN switch enables or disables a whole section, setting the SAT, COMP or LIM dial to 0 bypasses the according effect but not the whole section.

2.3. Some tips on using the presets

Rather then offering a variety of presets for specific mixing situations NastyVCS features some templates which you can use as a starting point and adjust to your needs. There are some “Preamp simulation” programs, some for “virtual summing” and some other dedicated to certain program material or mixing tasks.

Always adjust the audio input level (which “drives” certain components of the plug-in) and the output level (for A/B tests).

3 Advanced Usage

3.1. Understanding gain-staging in NastyVCS

The audio signal flow in NastyVCS is basically from left to right through all the components as shown in the interface (with few exceptions as discussed in Chapter 3.2).

Beside that, there is just one single concept to be understood to know how gain-staging is actually working in NastyVCS internally. This affects the SAT, COMP and the LIM section and works identically with each of them and so just one of them is explained here.

Gain-staging in each of these sections is basically working from the bottom to the top which means:

1. The IN switch determines if the whole section is in or not and therefore if any level changes are applied or not.
2. Next comes the lower of the two knobs which can alter the audio volume as a side effect (SAT in this case).
3. Note: If this knob is set to 0 this is a true bypass of the effect (but not of the whole section).
4. The last step is the knob above (IN in this case) which alters the audio level after the main effect (SAT in this case) has been applied.

The concept of a true bypass as hinted in step 3 allows to use one section for just gain-staging purposes and not applying any other effect. For example if you don't want to apply any compression effect but want to drive the signal upfront the limiter just use the compressor stage and set the COMP dial to 0. Compression is now entirely bypassed and you can simply use the MAKEUP dial to adjust the audio level.

3.2. Side-chaining and routing

Basically the audio signal in NastyVCS flows from left to right through all the components as shown in the interface. There are two exceptions for this, one is the sidechain signal flow and the other are the routing options.

The sidechain signal is basically the audio signal which is presented to the compressor to calculate it's gain reduction amount from. In NastyVCS this can be obtained from the main audio inputs (channel 1 and 2) but can be selected from channel 3 and 4 as an option by switching EXT SC (external sidechain) on. Your host must support this, of course.

In both cases the filtering stage (HP and LP) can be switched to appear in the sidechain path and not the main audio path anymore. This is done with the SC switch in the filter section.

By default, the EQ section is in front of the compressor. This can be altered by turning on POST in the EQ section. The EQ is then located between compressor and limiter.

3.3. Saturation and phase alignments

NastyVCS offers a dedicated input stage which allows simple input gain control but features crunchy “pre-amp” saturation as well. The amount can be dialed in to taste and allows very first peak treatments ranging from subtle to nasty (sic!).

Preamp stages usually are affecting both, saturation and phase alterations but in NastyVCS this is separated to allow way more and precise alignments of the audio signal's phase. The phase section is located left of the compressor and the combination with the compressor and/or the saturator then is the ticket towards “preamp style” or “opto comp” audio signal path coloring.

The input saturation features harmonic distortion and is a good alternative or complement to limiting. Example: on difficult to handle bass tracks the input stage can do wonders the limiter can’t. Finally, the combo of both can be a real killer in some mixing situations.

The phase tool basically combines all-pass filtering in a special setup plus some additional controls. The basic phase shifting shows up as in the diagram below (90/180 phase adjustment engaged). It features an additional polarity switch (also referred to as phase inversion, hence the INV switch label) and a phase frequency center knob. The diagram shows different PHASE knob positions ranging from 0 (right-most curve) to 12 (left-most curve).

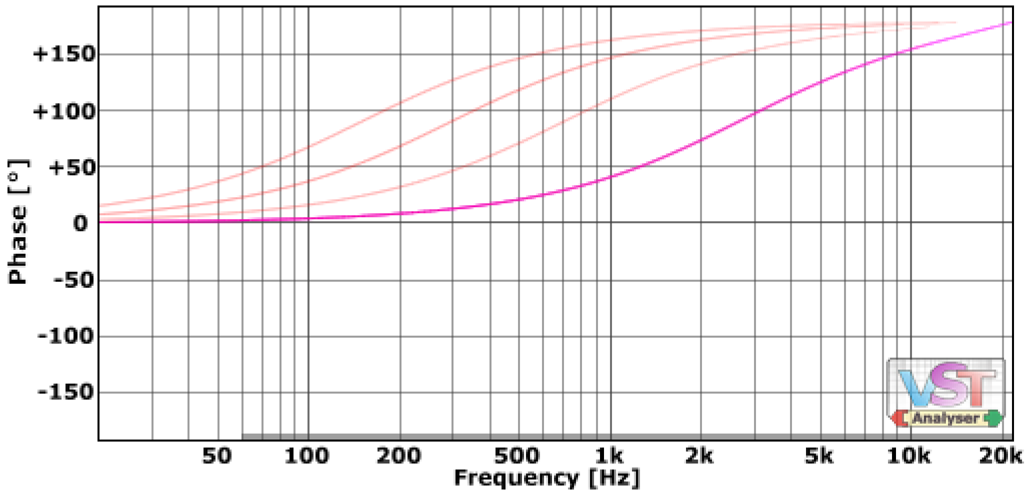

The interesting thing about this approach is that even if some rather constant phase shifting approaches (which are possible in the digital domain) might appear to be more reasonable or accurate in theory the method still maintains a more interesting or musical compensation at the end (since phase alteration in the analog domain does not appear as theoretical ideal).

Such alignments can be used for rather “artistic” mixing treatments and audio signal coloring as well. Just as a simple example, adjusting it properly to the upper frequency range one can easily achieve some serious amounts of transient smearing or, applied to the lower frequency part, the bass range can be decoupled to some extend (some bass enhancers are taking advantage of such effects as well).

3.4. Compression and limiting

The heart of NastyVCS is the entirely new build opto-electric style compression unit. This type of circuit designs do have a significant and highly program dependent behaviour and this is mainly due to adaptive release time characteristics, the typical compression transfer curves and the overall frequency dependency and non-linearities (have a look to chapter 4 for more details). Result: smoothness even on difficult to handle audio material.

An external sidechain option is available for the compressor (engaged by EXT SC) and with SC turned on the filtering section can be routed into the sidechain. If you prefer the EQ being behind the compression stage then engage the POST switch in the EQ section. Last but not least, the sidechain signal can be boosted with the +12dB switch to easily increase the compression amount.

While creating this compressor, major efforts has gone into designing the specific attack and release time characteristics which are available as fixed programs: three programs for setting the attack and five for release time behavior. Those programs cover a really useful range from faster to slower timings but since everything acts program dependent no precise timing information is specified. They are sorted timing wise from fast (top) to slow (bottom).

This compressor maintains *punch* in almost every situation and you can't really go wrong with it – it's a stellar but easy to use mixing comp. If you are in the need for more explicit dynamics treatments with NastyVCS then simply choose the limiter. Placed right to the end of the signal path, the limiter gives the final transient control and can smoothly be dialed in from zero to 100%. Set to 100% it works as a true and accurate brickwall safety limiter to prevent peaks to go further above 0dBfs.

As with each and every other NastyVCS component, the limiter does not need any lookahead information and so even limiting is available for latency free tracking. Two timing options (selected by the FAST switch) can be chosen, complementing the compressor for final management of audio transient information.

The internally calculated amount of gain reduction can be displayed visually with the COMP and LIM option right below the vertical VU style meter. On top of that, the limiter features a little clipping indicator (small LED right beside the LIM label) which roughly indicates clipping events.

3.5. EQ and filter

Similar to the EQ's and filters in BootEQ mkII, NastyVCS stays on the musical (and not the surgical) side of the audio source. It is more a coloring toolbox (especially in the combination with the PHASE option) rather than allowing to shape audio beyond recognition. Some classic technical principles have been carefully selected and replicated and the overall combination and attention to detail makes NastyVCS stand out in the crowd of equalizers today. Given the specific curve selections you might notice its slight “old-school” attitude when mixing through NastyVCS.

To get an idea of this and the overall philosophy of NastyVCS let's have a closer look at some of the technical designs starting with the boost/cut style equalizer. Boost/cut style equalizers are typically obtained from EQ's working in parallel configurations and “taking advantage” of certain effects coming from their interacting phase response. In NastyVCS I did not only want to have a basic boost/cut EQ, I wanted to have the boost curve from one of my favorite bass EQ's, combined with the cut option to allow those classic boost/cut curves. This ended up in a way more complicated technical design which actually combines some serial and parallel filter configurations.

Likewise with the mid frequency bell type EQ's their boost and attenuation behavior is calculated individually, using different algorithms.
Following the same concept, the HP in the filtering section is implemented as a straightforward 12 or 24dB butterworth filter while the LP features a different and

smoother design which avoids warping near the Nyquist frequency and offers gentle high-end treatments.
Last but not least the special AiR shelving EQ resembles the top end curve of a specific mastering EQ but also adds that certain slight “dip” below the boost which is so well-known from another type of (musical sounding) mixing EQ.

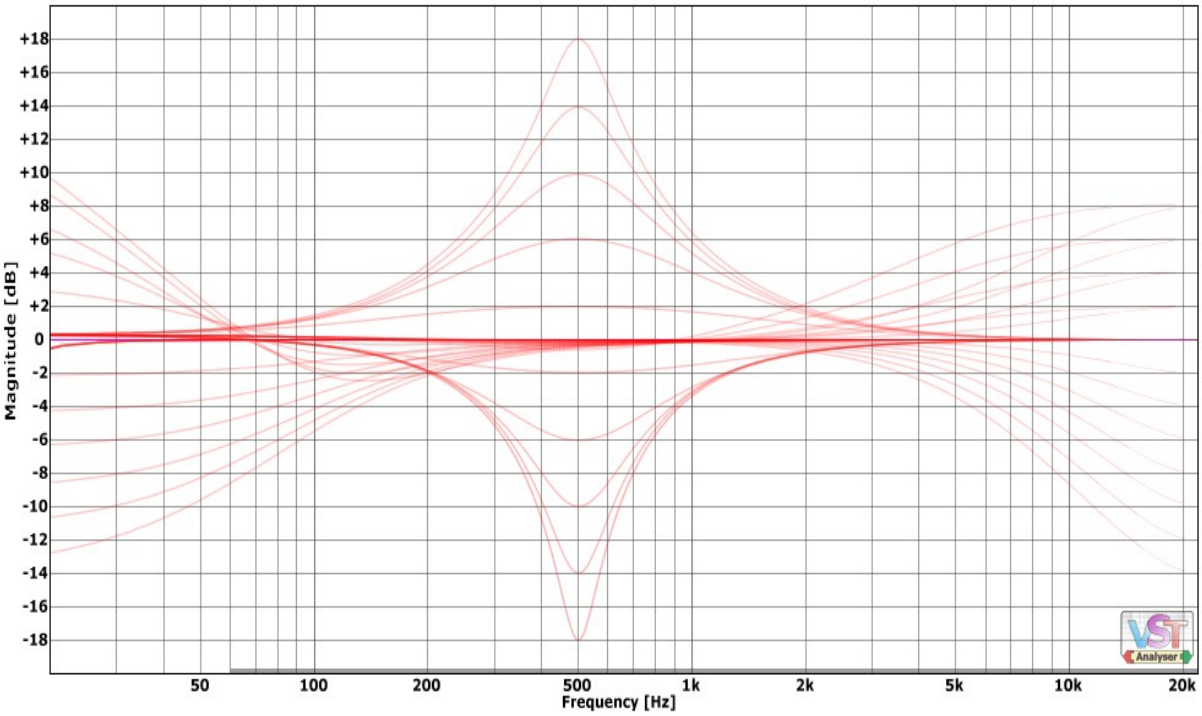

3.6. Putting it all together

NastyVCS really shines when all its tools are properly combined. Condense the dynamics in some small amounts right in the input stage, apply some smooth compression without affecting the punch/transients and then let the limiter eat some peaks at the output. Combine saturation and phase shifting to obtain some fancy audio signal coloring or just use the phase option for pure alignment tasks during recording. Gently limit the audio frequency range with the smooth filters or use them in the sidechain while adding some musical texture with the 'old-school' equalizer curves.

4 Addendum

4.1. The beauty of opto-electrical compression

Opposed to VCA, Variable-Mu or FET based approaches, opto-electrical compression takes advantage of using a light-sensitive resistor and a small light emitter (a LED or electroluminescent panel) to obtain a gain reduction voltage in the sidechain path. This technique is well-known to add some smoother gain riding characteristics to the signal because of the specific attack and release response which comes from the inertia and inherent memory effect of the photoresistor element.

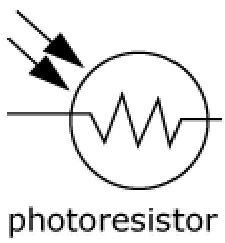

Program dependency

Typical classic opto-electrical circuit designs do have a significant and highly program dependent impact on the processed sound and this comes mainly from three factors:

1. The adaptive release time characteristic which gets faster at higher compression activity and can be fairly long when leaving the compressors duty cycle.
2. The specific compression transfer curve which features soft-knee characteristics by nature and limits the dynamic range.
3. The frequency dependency and non-linearities impacting the actual behaviour of things like the compression transfer curve.

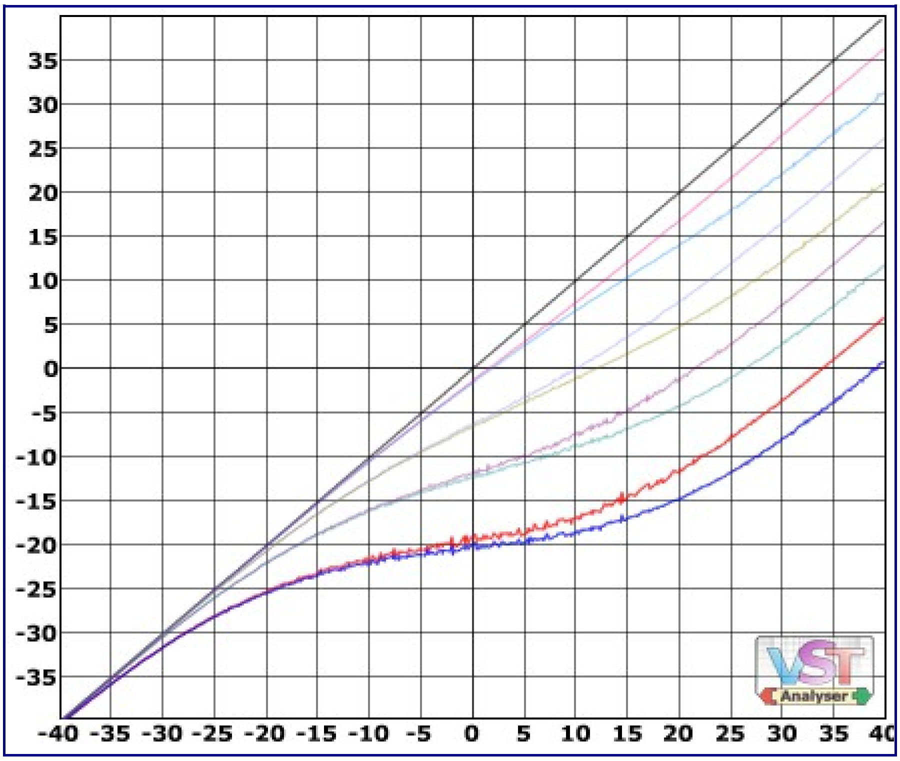

In addition, opto elements do have an inherent lag time in their attack response which is typically not fast enough to catch short transients but adds up to an overall smooth gain riding impression.

Light and shadow

Opto compression really shines on overall adaptive and smooth gain riding purposes such as for vocals and solo instrument performances, whereas faced with rather complex program material it can easily sound a bit quirky. Arguably that's why it is less ideal on the mix-bus, at least if you need bigger amounts of gain reduction.

In the digital domain the effects of the opto element are artificially modeled anyway, so these drawbacks can be avoided and both frequency dependency and non-linearities can be applied nicely to full program material.

4.2. About audio signal coloration

In this comprehensive article some deeper explorations and explanations on this topic are given and at the end a brief but handy definition about audio signal coloration is proposed. Some tips on mixing can be obtained here as well and – by the way – some myths about equalizing audio in the digital domain get busted.

But first let’s have a closer look into a different domain, the domain of digital image processing. In digital image processing, the fundamental color impression of an image is actually changed by performing some proper DSP maths on parts of the color spectrum of the image as shown in the example diagram above. Typically, a digital image is encoded into a 3- or 4-dimensional color space (like RGB or CMYK) and then each dimension can be manipulated individually over the spectrum. This changes then the overall coloration (and other things like brightness or contrast as well).

Quite similar, in the audio domain there are two dimensions over the frequency spectrum which can be utilized to alter the audio *color* impression: the magnitude and the phase response and this is typically done by an equalizer or filter. We are not going to talk here about those drastic phasing effects which are typically introduced by time shift based effects such as chorus and delay. While the impact of alterations in the frequency magnitude response curve is quiet obvious, altering the phase response might be not (and is often mixed up with other EQ side effects like resonance or ringing of a filter).

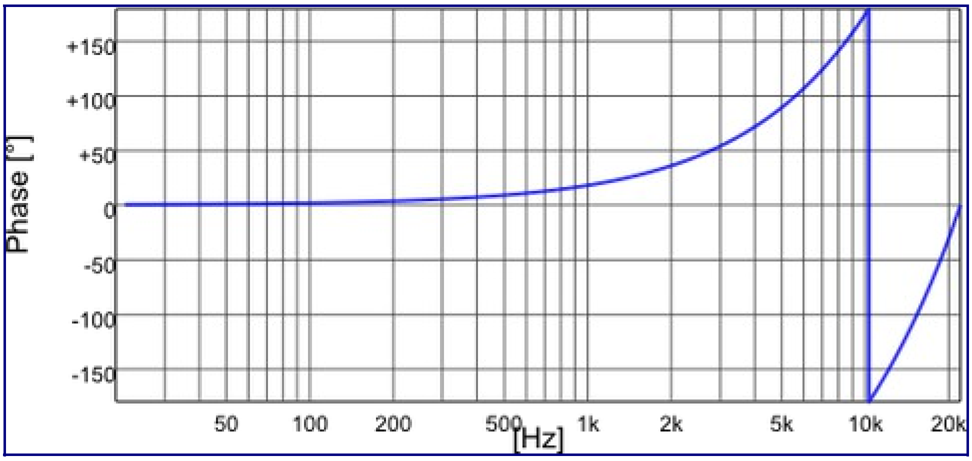

So, how does a certain phase alteration actually affect the sound perception of the audio? Of course this depends on the real frequency where the phase alteration (aka phasing, phase shift, phase drift, phase warping, phase distortion) occurs but for the sake of simplicity lets first have a look at the rather general effect caused by an overall and continuous phase shift: Lets assume we have an effect which introduces a continuous phase shift over the entire frequency range but changes nothing else (which can be performed by an Allpass filter – see the figure above). Now we have three scenarios, depending on the amount (degree) of the phase shift:

1. Slightest phase shift: Human ear does not perceive anything and thus can't judge any better or worse sonic quality.
2. Some amounts of phasing: The phasing, which causes some drift of higher frequencies in time, can now be perceived by our hearing. This slight displacement in time could be perceived as to be a more sound, less edgy and harsh audio quality and even to have more room / depth impression. (Hint: this are qualities some people hear and associate to the rather positive properties of analog audio processing).
3. Larger phase drifts: Larger displacements in time increase this effect and can completely destroy the transients. The signal is perceived as being washy and roomy now and lacks definition which is not desired in most cases (but can be

useful e.g. as part of reverberation processors).

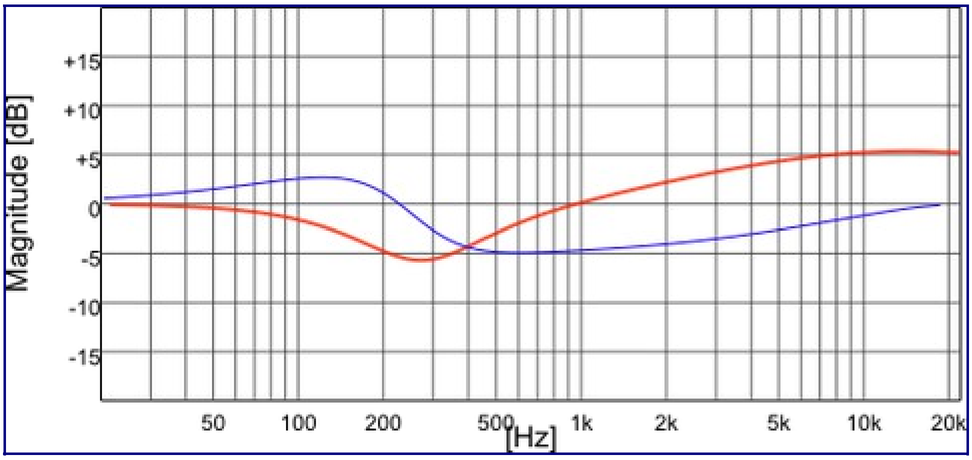

In digital reality, phase shift is not introduced that much over the entire spectrum by our commonly used DSP mixing effects such as phasing EQ's but has a rather local effect and is not that easily detected by ear since the phasing effect is being concealed by the frequency magnitude change of the EQ. This holds true at least as long as gentle and broad magnitude changes are performed. That said, if rather deep and steep changes are made then also larger and maybe unwanted spectrum displacements are introduced by such an EQ.

This may lead to some serious issues when for example on each and every track most signal resonances were removed by such steep filtering effects which is a common misconception in mixing audio. It leads not only to rather flat and boring signals but also introduces significant phasing issues as described above as long as no linear phase EQ is used (which introduces other problems and is not discussed here) and as an overall result the mix gets fluffy and lacks definition. As a side note, this also shows that the prejudice that cutting is always preferable to boosting other frequencies is an urban myth.

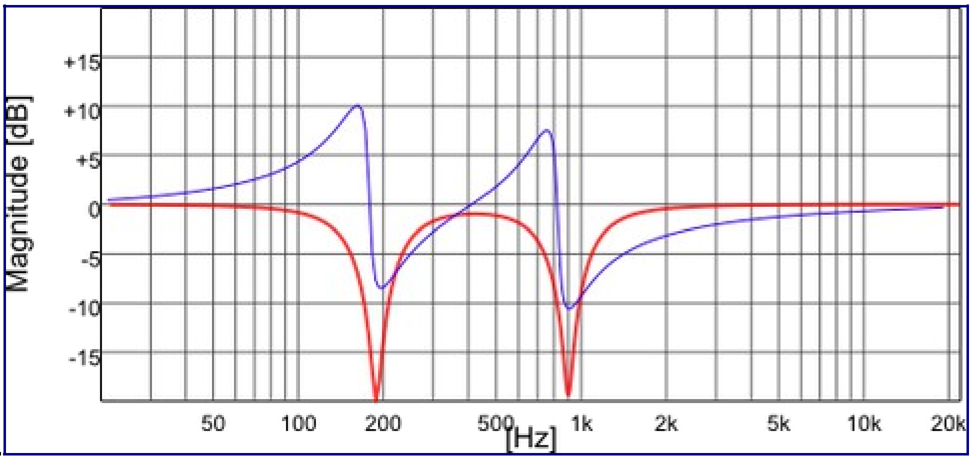

.

In some cases it might be considerably better to gently boost the desired frequencies instead of deeply cutting some unwanted ones and the simple "garbage in, garbage out” law applies here too: If that much and rather deep cutting or filtering in general would be necessary on such a signal then it’s probably better to try to fix this by changing the source, the recording situation or the arrangement. A good excercise is to set up some sound sources plus arrangement where (almost) no EQing is necessary during the mix. As an added sugar, such well-selected/recorded and arranged sound sources typically lead to a way better and much more natural loudness performance in the end. But back to topic.

So, is this phasing really bad and has to be avoided in any case? As hinted earlier, applied in slight doses, phasing can introduce a very nice and pleasant audio signal coloration and is part of the sound that we typically associate with high quality audio processing in the analog domain. In DSP land, these kinds of effects can be used to the audio engineer's advantage as well by simply applying phasing in the right amounts and in the right place of the spectrum. This is implemented e.g. in some audio enhancer circuits which are introducing dedicated phase shifts aimed at the spectrum or specific to the loudness performance of the audio signal.

Even some digital compressors are utilizing this. This answers yet another interesting question - can we also introduce audio signal coloration just by using plain dynamic effects? The short answer is: Yes we can! This is rather obvious and can easily proved by an audio analyzer in many,though not in all cases. Some compressors e.g. introduce gain reduction dependent phasing which can subtly change the color impression and of course a true multi-band compressor is able to perform drastic frequency magnitude changes for obvious reasons.

But even if the dynamic processor does not alter the frequency or phase response in a direct fashion it can alter the perceived spectrum just by having implemented a frequency dependent sidechain treating parts of the spectrum differently from others in respect to their loudness performance. This is also true for transient processors in general when transient information is typically associated to (and treated in) specific frequency ranges.

The phrase “audio signal coloration” could simply be seen and understood as “affecting the perceived tonal spectrum” no matter which phenomena or method actually caused it. A phasing EQ, when properly applied, is a good way to pleasantly color the audio in both dimensions, frequency magnitude as well as phase response. A linear phase EQ colors the audio too but just in one single dimension. Other processors such as compressors or enhancers can potentially take advantage of audio signal coloring as well, not even mentioning the time shift based effects such as chorus or delay.

4.3. Judging saturation effects

There are quite some misconceptions around on how to judge a saturator's sonic quality, here are some tips to avoid the most common pitfalls:

1. A good saturator does not appear as distortion in the very first place. Firstly it just saturates incoming audio signals which means that at a similar RMS output level it simply reduces the peak performance (which results in a smaller “crest factor”).

2. This immediately implies that you need a RMS meter in your output chain to compare different saturation settings or devices to another. Basically this is the same for comparing limiters or maximizers.

3. Distortion is a side-effect which typically occurs at higher saturation levels. It can have different sonic qualities, e. g. due to the frequency distribution of distortion which makes a huge difference to human hearing and whether the effect is perceived as gentle or not.

4. Don’t rely on a simple spectrum analyzer here, it does not know anything about the concept of being “gentle” or not.

Summary: Always assure equal RMS output levels and then use your ears.

4.4. Updates and further information

Refer to my Blog at http://varietyofsound.wordpress.com for some additional information and updates on this plug-in or leave a note there if any issues did occur.

Peace,
Herbert