Dynamics Toolbox Audio Processor

White Paper

Issue 1
November 2007

Safe Sound Audio, UK

Introduction

Safe Sound Audio's first product, the P1 analogue processor, has gained a strong reputation in music studios for its open sound and very smooth multi-sideband *peakride* compressor. The Dynamics Toolbox is a line level dual channel, stereo linkable version which preserves the purity and warmth of its predecessor and adds a whole set of new enhancements based on feedback from P1 users.

Design philosophy

Safe Sound's goal is to create a new generation of analogue processors which deal with some of the limitations of traditional analogue processing especially in audio compression and limiting.

Peakride compression, as introduced in the P1 audio processor, uses three linked sidechains each with different attack, ratio and release characteristics, which in tandem allow for very musical dynamic control of source material, especially on vocals and acoustic instruments which can often suffer badly from over compression in an effort to bring fast transients under control. *Peakride* compression evens out levels in a very musical way without robbing the sound of life, so bringing the performance a new energy whilst remaining transparent and not sounding overly effected.

Responding to user feedback

The three most asked for enhancements from users of the P1 were;

- A two channel stereo linkable version
- Access to the compressor sidechain
- SMOOTH compression is good but sometimes we want more user control.

The Dynamics Toolbox answers these requests and a lot more too. The two channels are fully linkable with channel 1 controlling all dynamics functions in 'link' mode, so there is no need to match left and right settings manually.

The sidechain is fully accessible to the user via a set of balanced insert jacks, but we also thought it would be useful to provide a set of fully sweepable hi and lo pass filters so dynamic EQ effects are possible without the need to patch in an external equaliser. The very popular 'big' bottom end sound is easy to achieve and fine tune either in mono or stereo when tracking or mastering.

The third request made us pause for thought. *Peakride* compression is very smooth and easy to set up, but its design does not lend itself to extending user control especially in regard to manual setting of release times.

So we went back to the drawing board and invented a brand new compression technique called 'dynamic tracking', and made it available as an alternative to *peakride* at the push of a button. Although, for convenience, they share a common set of user controls, the design philosophies are quite unique.

We also had a lot of requests for the P1 limiter to have an adjustable threshold so it could be used as an effect in addition to its usual work role of protecting against digital clipping when feeding into DAW's. So we redesigned the limiter in the Dynamics Toolbox to meet the request BUT then came an unexpected twist in the tale when we discovered that some recording engineers were using the P1's limiter in parallel compression applications.

New York Compression

'New York compression' became a trademark sound of many New York based mixing engineers and uses a form of parallel compression which mixes in a small amount of very heavily compressed stereo submix (usually drums or rhythm guitars) back into the main uncompressed mix. The beauty of this technique is that since you have full control of the blend between compressed and uncompressed submix strands, you retain the timbre and dynamics of the original submix whilst adding gritty texture via the heavily compressed element.

> From Wikipedia,
>
> *'Parallel compression, also known as New York compression, is a technique used in sound recording and mixing. Parallel compression is achieved by mixing a dry or slightly compressed signal with a heavily compressed identical signal. It is thought to maintain the subtleties of a performance while stabilizing the dynamics. Additionally, the attack and release settings of the compressor can be set in a manner that causes the signal to "pump" or "breathe" in tempo with the song, adding an identifiable character to the sound. It is most usually used on drum busses.*
>
> *It is sometimes erroneously referred to as sidechain compression.'*

We thought it would be nice to offer this feature 'right out of the box' rather than having to set up separate aux busses in your mixing console or DAW. The Dynamics Toolbox provides not only the blend control between dry (uncompressed) and wet (compressed) audio but also optional EQ'ing of the compressor sidechain, and it can also bring the limiter into play as part of the compressed audio contribution to the overall effect.

New York compression can be achieved entirely within the Dynamics Toolbox in both mono and stereo modes of operation. Parallel compression also has other uses particularly as a subtle form of 'bottom-up' compression widely used in classical music.

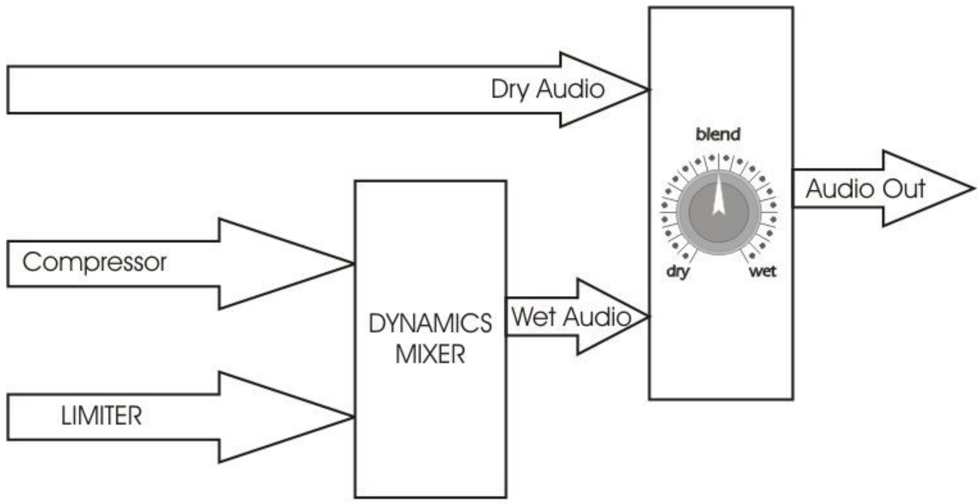

New York Compression
as implemented in the Dynamics Toolbox
DIAG 1

Another trick in the box

We were aware that many professional recording engineers swear by a particular brand of audio transformer. The Dynamics Toolbox offers a range of plug-in output transformers so that users can choose a brand which works for them. So we're engineering options for two European, Lundahl and Sowter, and two US, Jenson and Cinemag; and we'll continue to add further options if the demand is there.

In summary we ended up with the feature set as shown below.

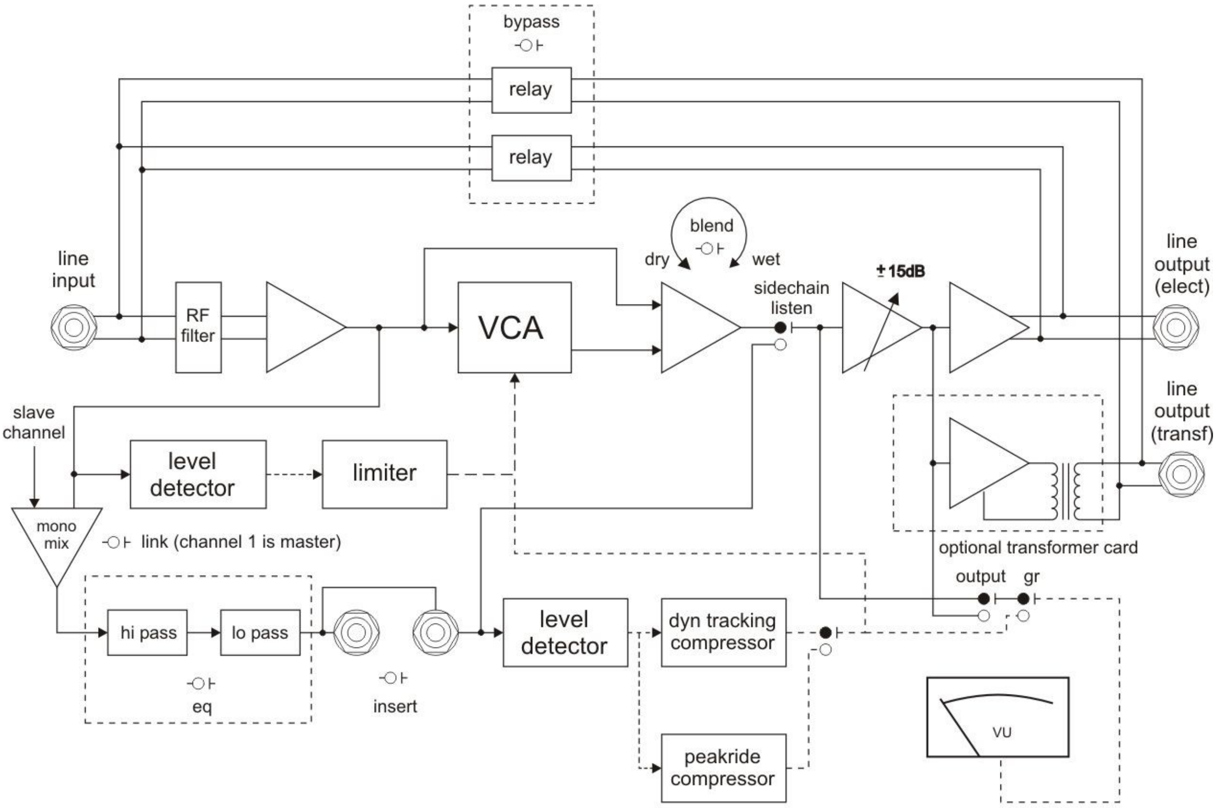

Dynamics Toolbox Audio Schematic
Diag 2

As is our way; here is chapter and verse on every aspect of the Dynamics Toolbox design;

The power supply

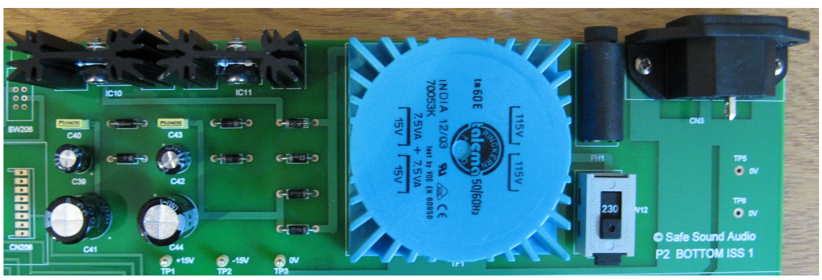

How many times have you cursed at the stupid plug in voltage 'adapters' so beloved by many of our competitors. Why do they use them? They're cheap that's why!

Most of these adapters and even many of the internal mains power supplies in audio processors are based on a principle called switched mode voltage regulation. It's an (almost) digital method of voltage production using a high frequency sampling wave, typically around 150 kHz. So your lovely analogue processor gets stuffed full of high frequency rubbish which is never fully filtered out. Not a very promising start!

The Dynamics Toolbox audio processor uses a fully regulated LINEAR supply with a very high quality toroidal mains transformer feeding separate voltage regulators. These provide ± 15 volts to power the 100% analogue circuitry within the unit.

Linear power supplies are less power efficient and more expensive to produce but they are the best way to power high quality analogue audio equipment.

LINE INPUT STAGE

The main audio path line input stages are electronically balanced using what we believe is the best device on the market for this application.

> From the THAT Corporation website,
>
> 'Developed by Bill Whitlock of Jensen Transformers, the patented InGenius input stage uses clever bootstrapping to raise its common-mode input impedance into the megohm range without the noise penalty from the obvious solution of using high-valued resistors. Like transformers, InGenius line receivers maintain their high CMRR over a wide range of source impedance imbalances — even when fed from single-ended sources. But unlike transformers, these wide bandwidth solid state devices offer dc-coupling, low distortion, and transparent sound in a small package at reasonable cost.'

The THAT Corporation InGenius® 1200-Series has proven ability to reject common mode noise including hum pick-up in pro-audio studio applications and provides such a strong audio performance so as to be near invisible in the main audio path.

In addition we've fitted RF filtering to all primary and secondary input stages to ensure good immunity from the ever increasingly crowded radio and mobile phone spectrum.

At this point a word on dynamic range and noise floor is appropriate. We designed the Dynamics Toolbox to have plenty of headroom to ensure it never clips even really hot audio but still with a very low noise floor so you can track, mix or master in confidence that your audio quality is being maintained. However different applications require different optimisations so we've designed in the facility to monitor at different nominal levels. Have a look at the two options below;

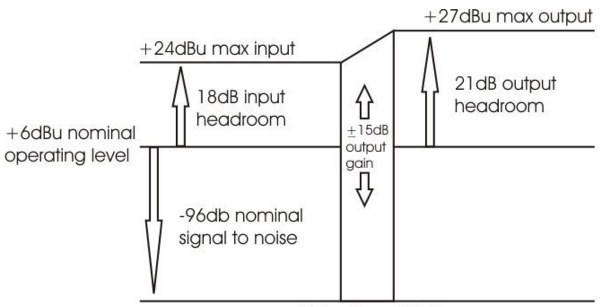

Level map : operating at +6dBu nominal operating level
TRACKING MODE
(DIAG 3)

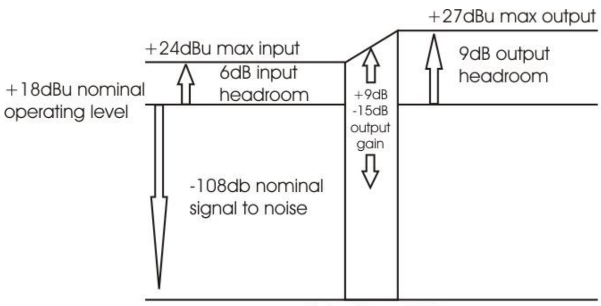

Level map : operating at +18dBu nominal operating level
MIXDOWN and MASTERING MODE
(DIAG 4)

When tracking it can be advantageous to work with at a lower nominal operating level. What does this mean exactly? Simply that you are setting aside some additional reserve headroom to deal with unexpected large jumps in audio level which can happen during tracking, so avoiding clipping, and ruining the perfect take. In the Dynamics Toolbox we suggest tracking at the nominal level of +6dBu (so +6dBu equates to the 0VU needle position on the DT's VU meters) which provides a whopping 18dB of input headroom, even more at the output stage, as shown in Diag 3 above.

During mixdown and mastering such headroom is generally not required so we have the option to switch to a higher nominal operating level of +18dBu (so now +18dBu reads as 0VU on the DT's VU meters) as shown in Diag 4 above. Now you are operating with the optimum signal to noise ratio but still with enough headroom to deal with the occasional over.

This option to choose your optimum operating level affects ONLY how you choose to monitor audio levels, it does not gain shift your audio signal at all.

THAT 4301 VCA

The heart of the dynamics gain control stage is a combined log-responding RMS-level sensor and voltage controlled amplifier chip also manufactured by the audio IC specialist THAT Corporation in the USA. This IC continues to be our favourite method of dynamic level control and combines a wide dynamic range with low distortion (under 0.01% at unity gain) in an easy to configure package. Also used at the heart of our P1 audio processor.

THE COMPRESSOR : FEATURING 'PEAKRIDE' AND 'DYNAMIC TRACKING' MODES

As we mentioned earlier, the *peakride* compressor first introduced in the P1 has been very popular especially for tracking vocal and acoustic instruments. So let's review the *peakride* mode of the Dynamics Toolbox first before moving on to the alternative *dynamic tracking* mode.

PEAKRIDE COMPRESSION

The original design goal of the *peakride* compressor was to have enough speed to catch those fast attack vocals but not to strangle the life out of them. We saw so many comments from users of existing voice compressors that they struggled to bring vocal transients under control, some even resorting to using limiters in series with compressors (actually not a bad idea!), only to complain that their super fast compressors were strangling the dynamics of the vocal through over-compression!

Experimentation with single side chain compressor designs led to three important findings;

- Fast attack compressors were very desirable for many vocal types but often led to over compression of the vocal.

- It was quite difficult and very time consuming trying to set optimal attack and release times for many types of audio sources.

- Very long release times were not appropriate for many types of source material including vocals but were often used to disguise poor ripple distortion performance.

Have a look at a typical 'first phrase' vocal waveform below;

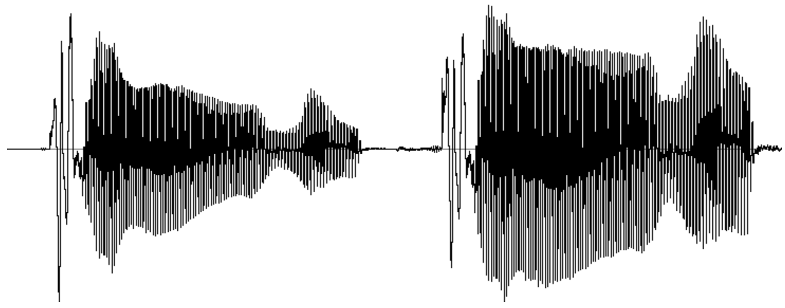

before compression after compression (with gain make-up)

Effects of over compression with 'short' 100ms release time
DIAG 5

The attack time of the initial vocal syllable is very quick so a fast compressor attack time will be necessary to bring some dynamic control to the vocal. If done with a simple fast attack compressor, this is what happens;

The first syllable attack is well controlled, but even with a short 100ms release time, the rest of the vocal phrase has been over compressed and most of the dynamic expression of the vocal phrase has been lost. It's a common misconception that percussion instruments have the monopoly on fast attack waveforms; vocal attack times can be staggeringly quick as well.

We tried shortening the release time down as far as 20ms, but now you start to get severe compression pumping, so that the compressor is gain pumping up and down with every vocal syllable. No good at all!

What's really required is a control chain which allows the compressor to 'ride' the crests of the audio so that reaction to fast audio peaks is fast and release from compression matches the natural 'smooth' decay of the audio waveform. It has to be capable of fast attack compression without the subsequent over compression shown above.

So when one control side chain won't do the job; use three!

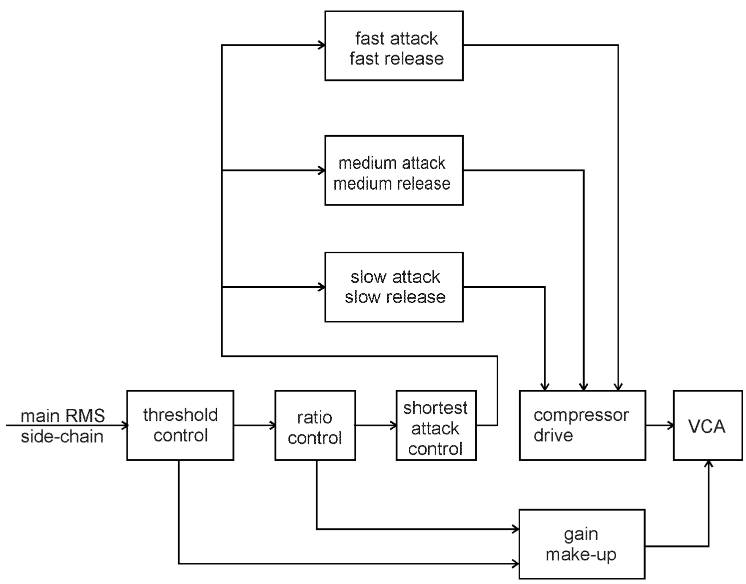

peakride multiple sidechains

DIAG 6

Side-chain 1 : a fast attack, fast release compressor which catches these over excited vocal phrases but recovers very quickly.

Side-chain 2 : a medium attack, medium release compressor which smoothes out the release response of side chain one.

Side-chain 3 : a slow attack, slow release compressor (but not too slow!) to provide very low distortion on sustained vocal phrases which often cause problems due to 'ripple' of the vocal modulating the control chain.

The three side-chains have different ratios and knee characteristics and are mixed together in varying levels of contribution to work in harmony providing a *peakride* compressor response.

The result is a variable (slightly) soft knee adaptive compressor which provides the following key advantages;

- Fast short vocal transients are well controlled without over compression.
- A programme related auto variable release time was possible making set up of the compressor much faster and simpler.
- 'Ripple' audio distortion has been almost totally eliminated without very long release times being necessary (longest release time is only 500ms).

Dynamic tracking compression

So whilst the *peakride* compressor offers a very smooth open sounding and easy to use compressor we had a lot of client feedback asking for a greater degree of user control of the compressor characteristics, especially manually set release times but without sacrificing smoothness.

Dynamic Tracking uses a novel technique to maintain low ripple distortion even at very fast attack and release times, so you get a very wide range of user settings but still maintaining a very smooth compressor.

Dynamic attack and release control techniques are quite different so let's deal with each in turn;

Dynamic attack

We had learned from our P1 design experience that fast attack compressors can suffer from low level distortion due to very rapid level control shifts. This becomes audible if fast attacks are maintained for long periods during compression (in this discussion we're talking about in excess of about 10ms, not seconds!).

Peakride gets around this problem by introducing a second (and then a third)

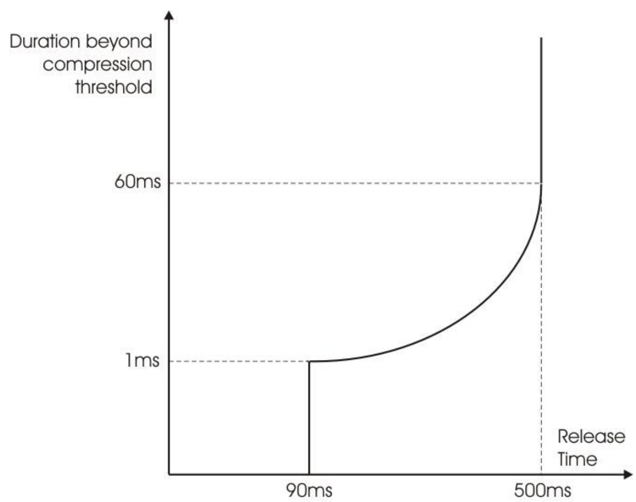

Peakride's signal dependent release time (DIAG 7)

more dominant slower attack cycle which takes over when the short burst fast attack has done its best work.

However the compromise is the introduction of longer release times as the audio signal spends more time above the compressor threshold, as shown in Diag 7 above.

This is why manually set release times are not possible in *peakride* compression.

So we wanted to achieve short burst fast attack times which slow down as the compression cycle continues without affecting our release time options. We did what all good designers do; we borrowed a technique from another of our products; from the P1 limiter design, as shown in Diag 8 below.

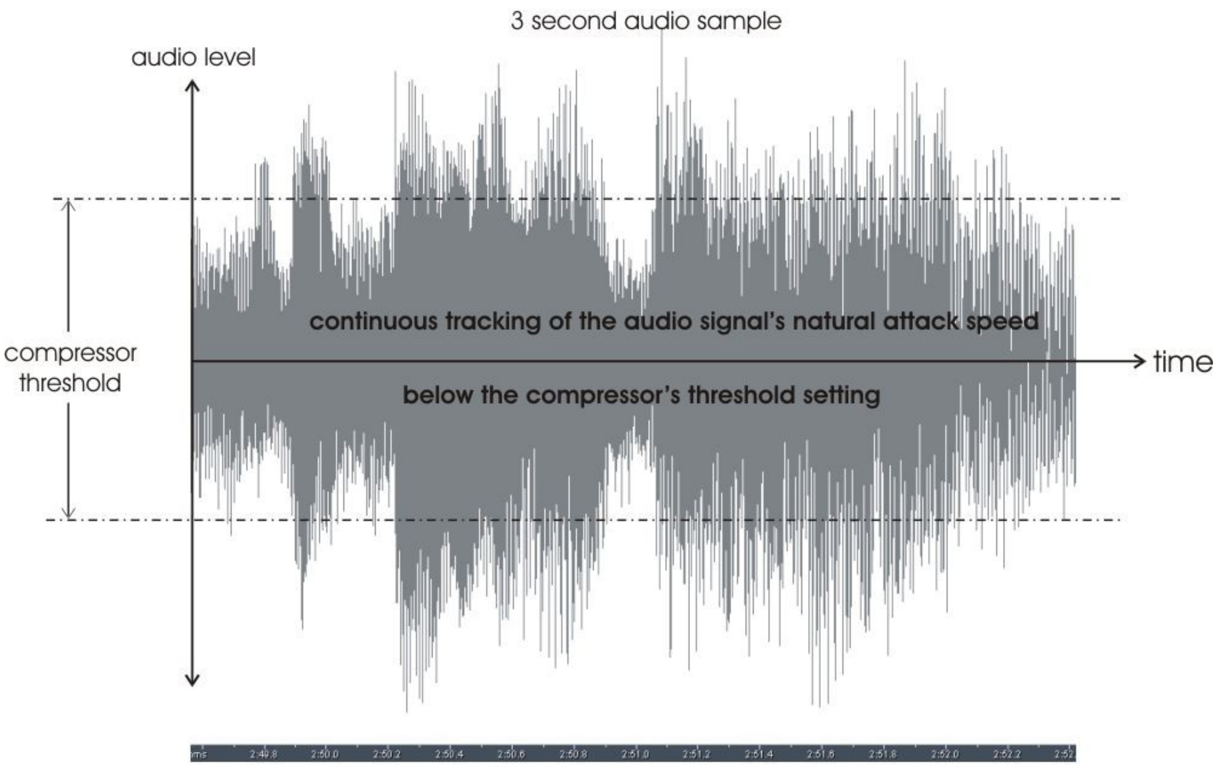

Tracking of audio signal's natural attack speed
below the compressor's threshold level
(DIAG 8)

Dynamic attack allows the compressor's attack time to track the natural attack speed of the audio signal. It does this by tracking the audio's natural attack speed before it reaches the compressor threshold level.

It uses this information to boost the compressor's natural attack time to match the needs of fast incoming audio but then dynamically reduces the compressor attack time once the initial audio excursion over the compressor's threshold has been dealt with.

The compressor's attack time continues to be boosted and then reduced through to whole period in which the audio signal's level exceeds the compressor threshold as shown below in Diag 9.

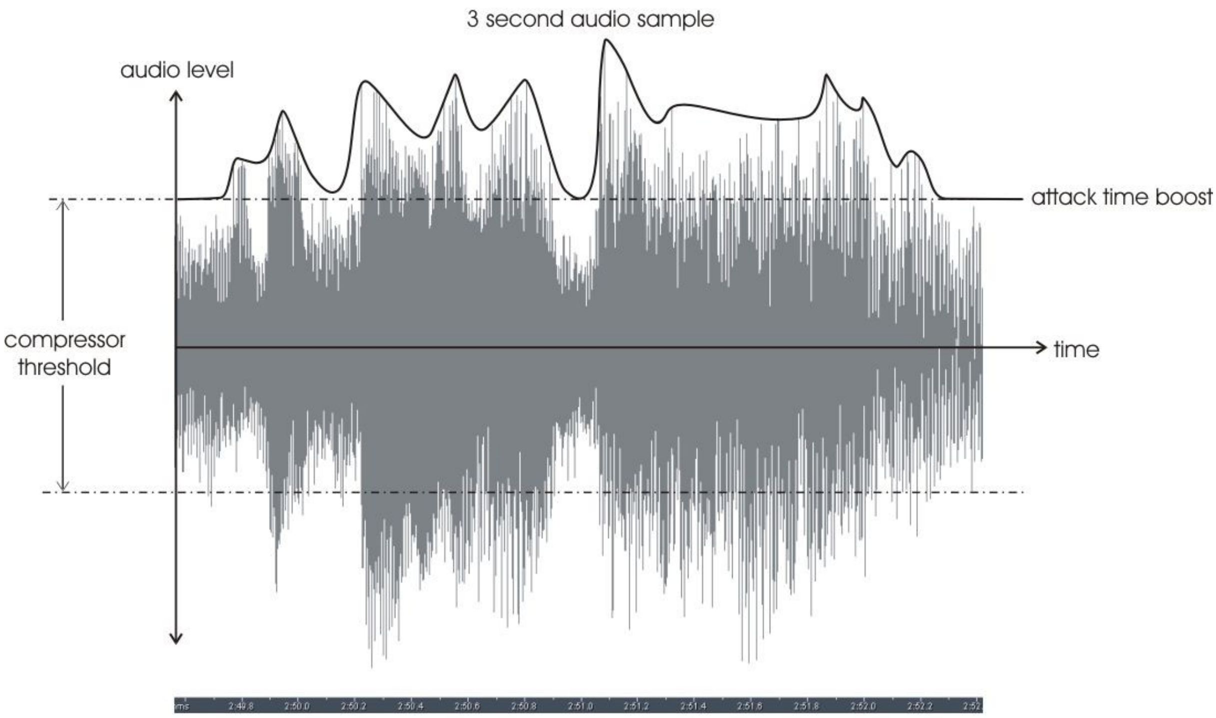

Compressor attack time boost mapped over the audio sample (DIAG 9)

What's the benefit?

In technical terms; lower levels of Fourier distortion (caused by very fast gain changes)

In practical terms ; it offers fast compressor attack times but without sacrificing smoothness.

Dynamic release

In 'traditional' compressor designs very short release times come at the cost of high levels of ripple distortion, especially at low frequencies as shown below in Diag 10 below. This type of distortion has a very negative impact on low end smoothness and can be especially damaging when trying to compress bass heavy mixes.

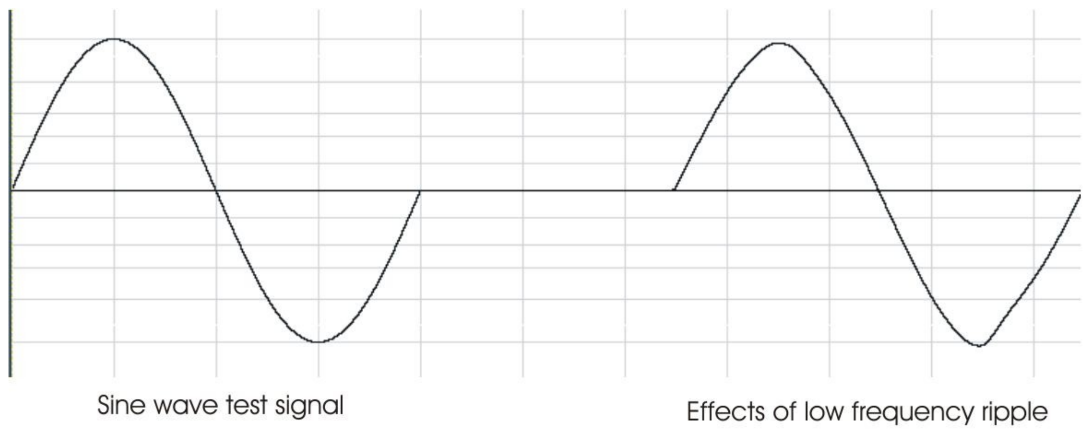

Ripple distortion due to very short compressor release times (DIAG 10)

Notice that the pure tone audio signal is beginning to lose its smooth top and bottom curves and begins to resemble a triangular waveform. This is not just an issue for bass guitar and bass drum compression. Many audio sources including piano and vocals have either a primary low frequency ‘carrier’ component and vibrato/tremolo induced low frequency components. All can suffer from audible distortion when being compressed or limited.

However the dynamic attack technique brought us an unexpected bonus when we started to look at the potential problems in offering very short manually settable release times in the order of 50 to 75ms. As the attack time begins to lengthen through the compression cycle (and remember we’re still talking about around 50ms audio sample times, not seconds) the attack boost applied is reduced and the dynamic effect is that the level of compression being applied to the audio begins to fall.

So even at very short release times the effect of slowing down the compressor attack time through each compression cycle is to smooth the transition between attack and release cycles.

The practical impact of this technique is a huge reduction in ripple distortion.

The *dynamic tracking* design is entirely compatible with *peakride’s* programme related release so we’ve added this as an ‘auto release’ option.

The summary schematic for the *dynamic tracking* element of the compressor is shown below. We believe it is a novel way to push the envelope of compressor design which allows for a very wide range of user settings whilst retaining the smoothness of response which is very much a Safe Sound audio signature.

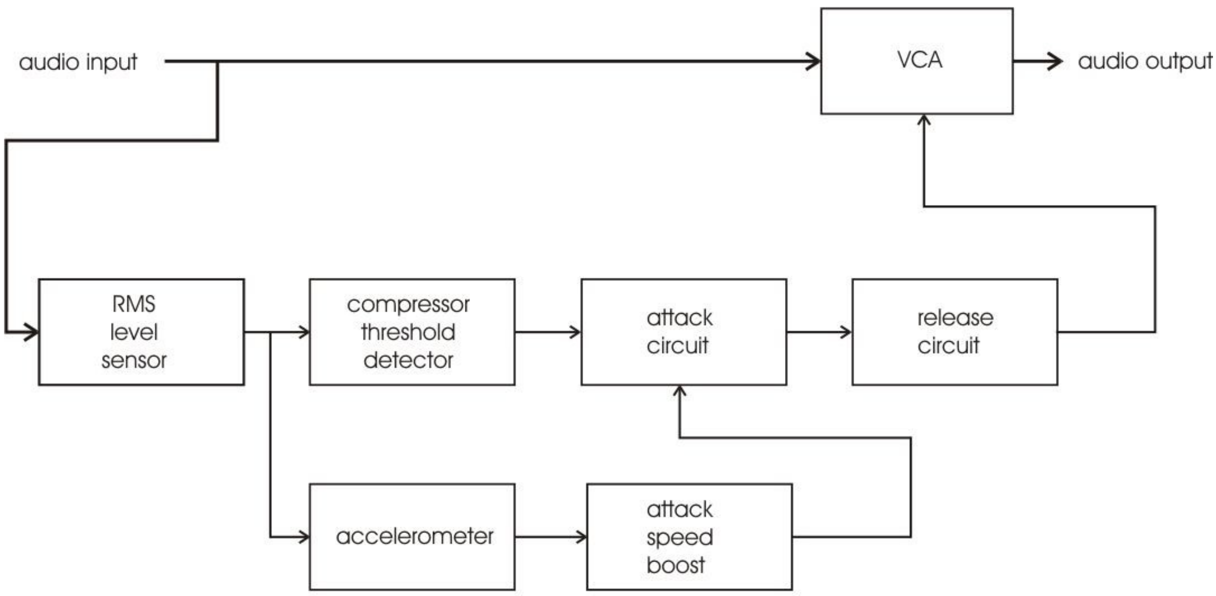

dynamic tracking compressor schematic
DIAG 11

THE LIMITER : FEATURING PREDICTIVE ATTACK AND PROGRAMME BASED RELEASE

For those of you who have browsed the P1 White Paper, you'll know that we went on a bit of a mission when designing the P1 limiter in order to offer brick wall limiting in an all analogue design. We got pretty close to that goal; so much so that Sound On Sound reviewer Paul White commented, 'Used anywhere close to sensibly, the limiter is effective and as transparent as any analogue limiter I've heard.'

Let's have a quick review of the issues.

True 'brick wall' limiting in the analogue domain is very difficult to achieve whilst maintaining high fidelity without the use of audio delay in the main signal chain. Have a look the Diag 12 below;

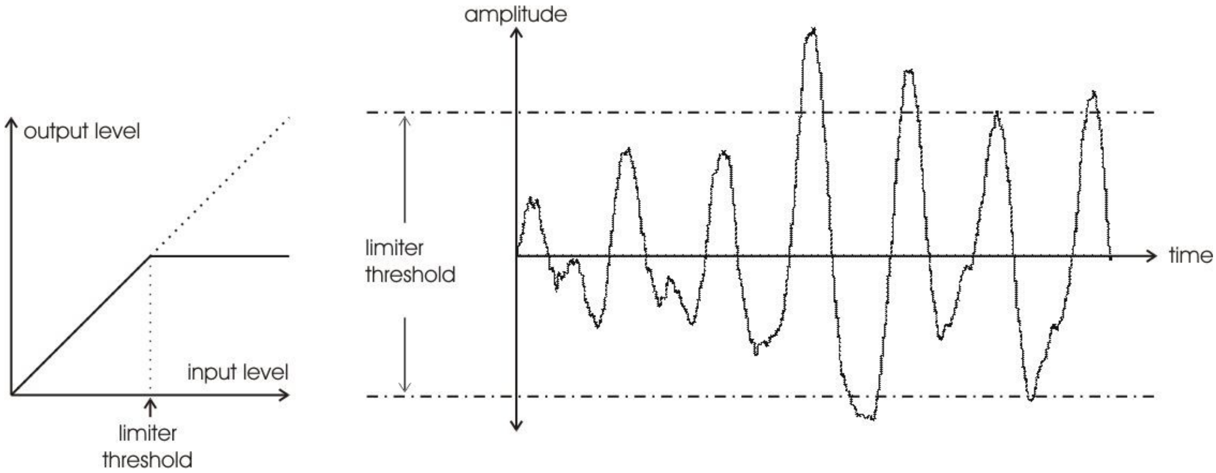

Basic limiter operation
DIAG 12

Audio is a time domain signal, measured as amplitude against time. If we consider a nominal limiter threshold; in the absence of an audio delay in the main signal chain, there is literally no time interval between an audio signal crossing the limiter threshold and, by definition, going into overload. This means that any mechanism used to limit the signal level has no time to react and so the audio will overshoot the limiter threshold until the limiter reacts and reduces the signal gain to bring the audio within the set limiter threshold.

If we added a time delay in the main audio chain, this would give us some time to 'calculate' a best fit gain reduction 'curve' but there are significant downsides;

- The limiter would be unusable in live performance.
- The limiter would be next to useless in live recording, making it impossible to generate a zero latency monitor feed for the performer. Even small delays of 2 or 3 milliseconds can cause problems, especially to vocalists. Look how hard soundcard and sound recording software manufacturers have tried to reduce 'latency' time for exactly the same reason.
- It would require the use of digital audio and a digital control chain which does not offer continuous control of an analogue signal and well……. it wouldn't be analogue.

The P1 limiter design got around these issues by combining three limiter sidechains all working with predictive dynamic thresholds which, just like our *dynamic tracking* compressor, track the natural attack speed of the audio waveform, in the case of the P1's limiter from 6dB below where it crosses the threshold level, as shown below in Diag 13.

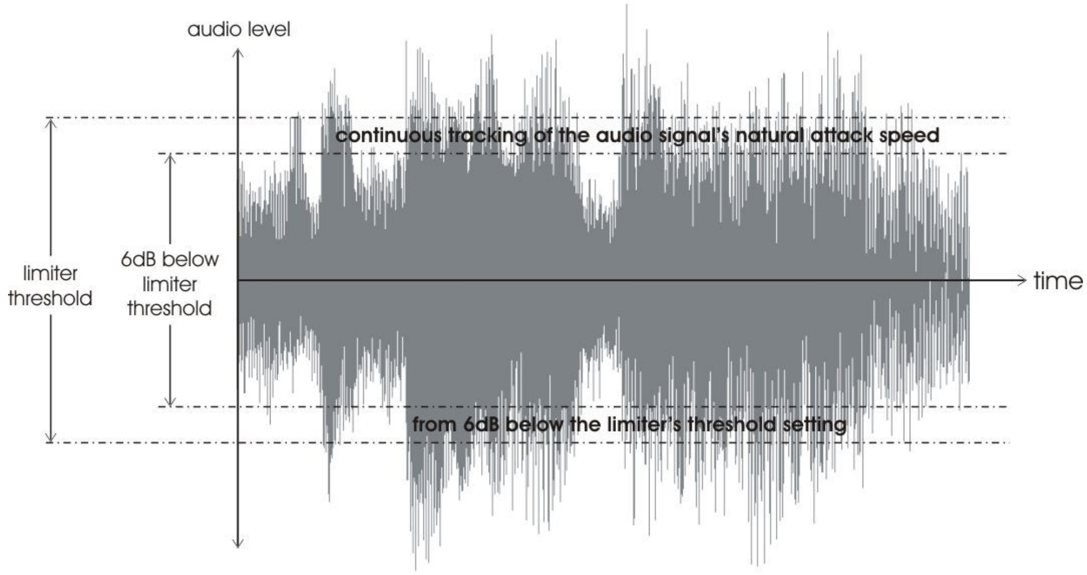

Tracking of audio signal's natural attack speed
from 6dB below the limiter's threshold level as used in the P1 limiter
(DIAG 13)

Adjusting the threshold dynamically is a technique which works really well, but in the P1 design it relies on a fixed limiter threshold so that the three sidechains track accurately. This is okay when using the limiter solely for protecting against digital clipping when tracking audio, but it was clear from user feedback that we needed to offer a variable threshold limiter within the Dynamics Toolbox.

We decided to retain one of the three P1 limiter sidechains within the new design so that small amounts of dynamic threshold adjustment could be made. This allows a degree of look ahead limiting but doesn't have any tracking issues. To this we added dynamic control of attack time, as used in the *dynamic tracking* compressor, but with much higher levels of attack time boost. This allows very fast short duration attack times to be fed into the limiter sidechain. The combined scheme is shown below in Diag 14.

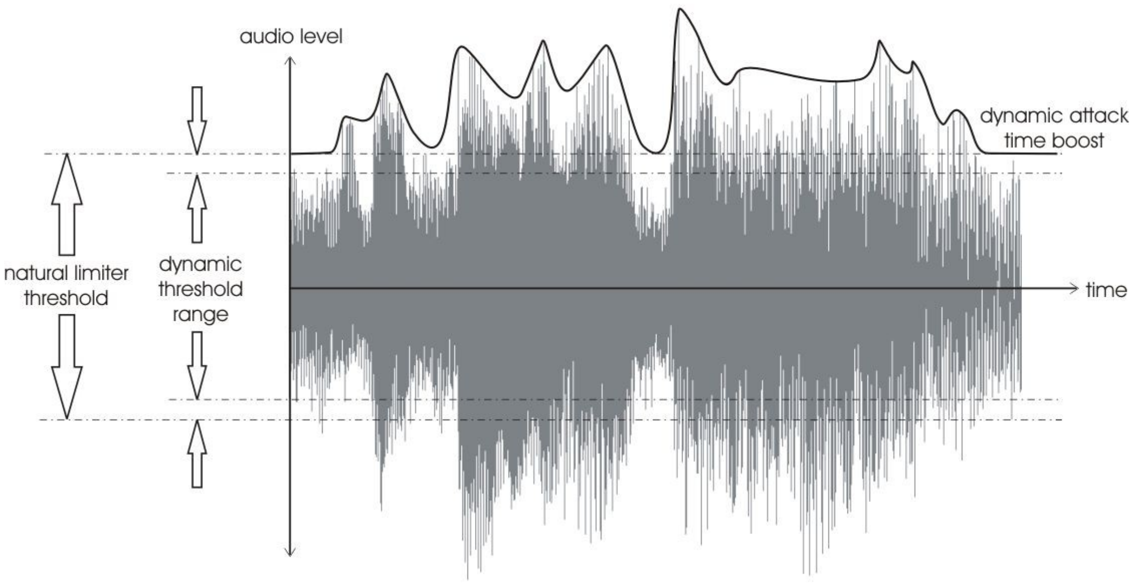

Combined dynamic threshold and dynamic attack time control
(DIAG 14)

The outcome is a limiter with burst attack times as fast as 100us but with the excellent low levels of ripple distortion usually only found in much slower attack time devices.

We've added a programme tracking auto-release circuit which provides for very smooth recovery from limiting adapting to the natural dynamics of the audio source.

The Dynamics Toolbox limiter design is shown below in its final form.

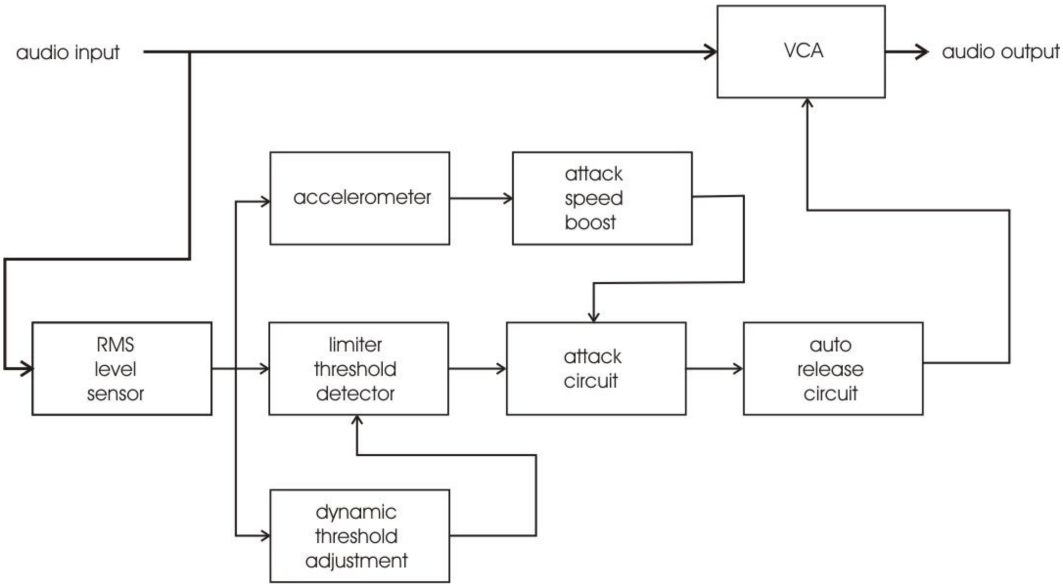

dynamic tracking limiter schematic
DIAG 15

Line output stages

Once we had made the decision to offer a number of optional plug-in output transformers within the Dynamics Toolbox design, which we thought would be very popular with many hi-end users, there might have been a temptation to skimp on the 'fitted as standard' electronically balanced output stage. We had the P1 output stage design to hand which offers a respectable +21dBu maximum output using an impedance balanced output configuration. But with many recording systems now offering +24dBu audio handling capability we decided to go one better and upgrade the Dynamics Toolbox audio output stage to handle +27dBu signals, so typically 3dB of additional headroom above most pro-studio set ups.

Once again we have opted to take one of the latest ICs from the THAT Corporation's 1600 series of audio line drivers. As configured within our design it offers superb low noise and low distortion and is extremely stable into long cable runs. It can also be connected into unbalanced loads with little loss of performance.

Plugging in one of the optional classic output transformers from Lundahl, Sowter, Jenson or Cinemag, whilst retaining the same +27dBu high signal handling capability, opens up a whole additional palette of tonal shaping for the professional audio engineer.

In summary

We hope you'll agree that the Dynamics Toolbox offers a very flexible sound shaping tool for tracking, mix effects and mastering. The number of combinations available through two compressor types, onboard equalisation of the compressor sidechain, and parallel compression techniques, provides an accessible and powerful toolbox of dynamics processing.

And remember that it's all accomplished with 100% analogue circuitry. In these days of data compressed digital audio formats we feel there are more reasons than ever to make high quality analogue processing available to professional sound producers.

As always we welcome user feedback to guide us in our future designs and look forward to hearing from you.