P1 Audio Processor

White Paper

May 2003

Contents

- **Introduction – Page 3**
- **The Power Supply – Page 3**
- **The Input Stage – Page 3**
- **THAT 4301 VCA – Page 4**
- **The Expander – Page 4**
- **The Compressor : featuring *Peakride – Page 5***
- **The Limiter : featuring multi-stage dynamic threshold control – Page 7**

Introduction

The aim of the P1 audio processor design is to create the first in a new generation of analogue processors which aims to deal with some of the limitations of traditional analogue processing especially in audio compression and limiting.

Much hype and mystic has grown up over the years about favoured compression and limiting techniques which claim to add magical properties to recorded audio. Sometimes this has some founding in fact but too often it is a clever way to explain away the limitations of the design. Turning a weakness into a marketing advantage is a great way to sell more product but doesn't necessarily help the sound!

Another common marketing ploy of many compressor/limiter designers is to tell the buyer just enough to wet the appetite without really explaining how it works. We decided early on to tell you exactly how the P1 achieves its excellent performance so here it is, chapter and verse;

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 150kHz. So your lovely analogue processor gets stuffed full of high frequency rubbish which is never fully filtered out. Not a very promising start!

We use a fully regulated LINEAR supply with a very high quality toroidal mains transformer feeding three separate voltage regulators. These provide plus and minus 15 volts to power the analogue circuitry and a separate fully regulated +48 volts to power condenser microphones. Unlike many products the P1 can properly power all condenser mic types without running out of steam, something which many 'adapter' based designs fail to do.

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.

The Input Stage

Three separate input stages optimised for function;

The microphone input features the latest Burr-Brown INA217 chip giving a strong combination of low noise and low distortion. We did also consider the Analogue Devices SSM2019 which has a marginally better EIN at maximum gain but the INA217 has a superior distortion performance in the middle gain ranges which is important for critical amplification of high quality condenser microphones. The SSM2019 remains an excellent second source device. Switchable +48V phantom power comes from a high quality linear regulator capable of driving even the most power hungry condenser microphones. We've included a permanent microphone input subsonic filter to protect against inaudible high energy low frequencies.

A dedicated instrument input stage has been provided which is necessary to provide the high input impedance required for direct connection of electric guitars and basses. It is difficult to marry high input impedance with low noise, so we found the best solution was to use a FET based input amp (provides a high input impedance) with a fixed 10dB gain stage. This lifts the signal level high enough so that the output noise of the FET input amp is not a factor in the achievable noise performance of this stage.

Both the mic and instrument stages normal into the final balanced line input stage which uses the industry standard 5532 op amp providing a really excellent noise and distortion performance and also a good drive capability for the insert send point.

Although it is less common to normal signals through the input jacks, both the microphone and instrument inputs are amplified before being normalled through the line input jack to the final line amplification stage of the input section. The normalling contacts in ¼" jack sockets have a large contact area and can at least be cleaned without dismantling the unit.

The use of a common gain control for all three input stages provides an easy to use and compact front end to the P1 and all inputs are level sensed to provide LED indication for any audio approaching to within 3dB of clipping prior to the limiter stage.

A switchable low pass filter with a turnover frequency of 80Hz and a steep 18dB per octave slope completes the input stage which is normalled via an unbalanced insert send to the VCA stage.

THAT 4301 VCA

The heart of the gain control stage is a combined log-responding RMS-level sensor and voltage controlled amplifier chip manufactured by the audio IC specialist THAT Corporation in the USA. This chip has graced many an analogue processor and we chose it due to its repeatable high quality audio performance and compact footprint.

The Expander

The decision whether to include a noise gate/expander into the unit was a difficult one. Most commentators warn again using noise gates on source material as what is done, cannot be easily undone after recording. However there are some good reasons for including this stage;

- The P1 does have potential as an effects processor for recorded material.
- The quality of many software based noise gates is still questionable.
- Expanders are very useful for shaping percussion sounds at source.

Having decided to include this stage, the next decision was noise gate verses expander. Noise gates have some advantages in signal to noise separation but the attack and release times have no relation to the audio, so even using a noise gate with multiple control knobs, you can spend forever finding the ideal settings.

Expanders are generally more forgiving because the expander characteristics are always related to the dynamics of the audio material. This was especially important in the P1 design where we wanted a simple to use single knob design. The basic design is shown below.

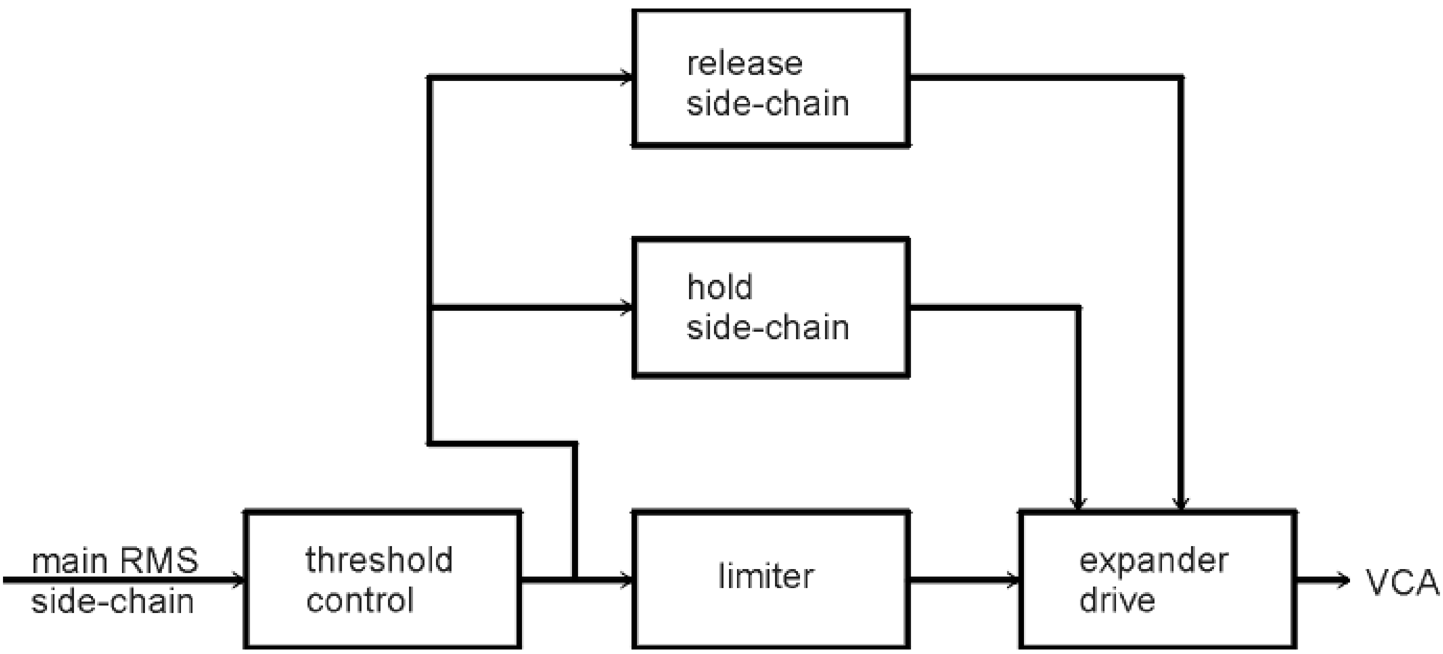

Although the *attack time* is nominally fixed at 3ms, this is the fastest attack time permitted in the design. If the source audio is rising at a slower rate then the P1 attack follows this. Attack times faster than 3ms are not a technical problem to design but have a tendency to produce clicks with very fast rising audio edges.

The use of an expander side-chain limiter is a method to reduce distortion when the expander is in circuit. It ensures that when audio level is above the threshold, there is no ripple distortion caused by the audio modulating the side-chain.

In addition to the main expander side-chain, the design uses two additional side-chains which in turn control the dynamic response of the expander drive stage, as follows;

The *release side-chain* follows the dynamics of the audio sample and measures how quickly the audio level falls when it drops below the expander threshold. This in turn affects the average release time of the expander.

The *hold side-chain* measures how long the audio has been above the expander threshold. This controls three functions;

- How quickly the expander goes into its release cycle.
- The knee shape of the release cycle.
- Also has an influence on the average release time.

These additional side-chains cause the expander action to compliment the natural dynamics of the audio signal.

Dynamic audio which bursts above the expander threshold for short periods gets the shortest hold and release times with a very hard knee.

At the other extreme, less dynamic audio which sits above the expander threshold for longer periods gets the longest hold and release times and a soft knee transition.

The fairly high *expander ratio* of 1:3 reflects the main intended uses of the P1 expander for noise reduction and dynamic shaping purposes.

The maximum *expander depth* of 20dB allows a fairly slow attack time to be used (which improves audio quality) but provides enough depth to have a significant effect on noisy source material.

The Compressor : featuring *Peakride*

Compressors with character, now there's a common catchphrase! but what about the audio quality?

The design goal of the P1 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. 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;

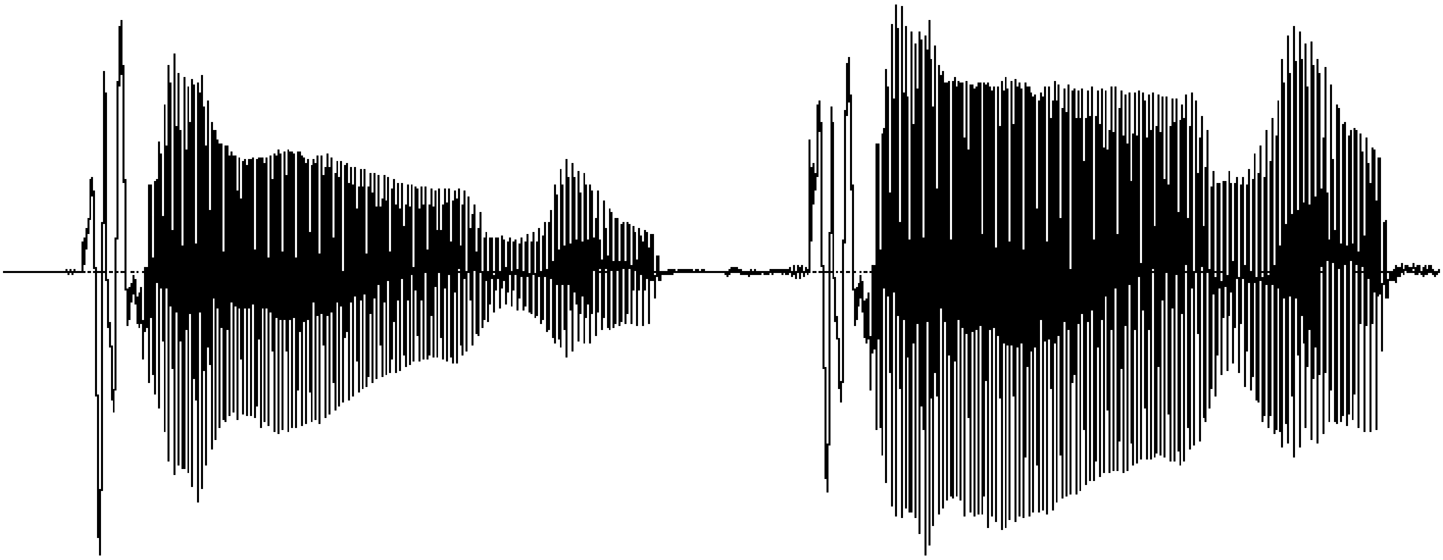

Before compression After compression (with gain make-up)

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!

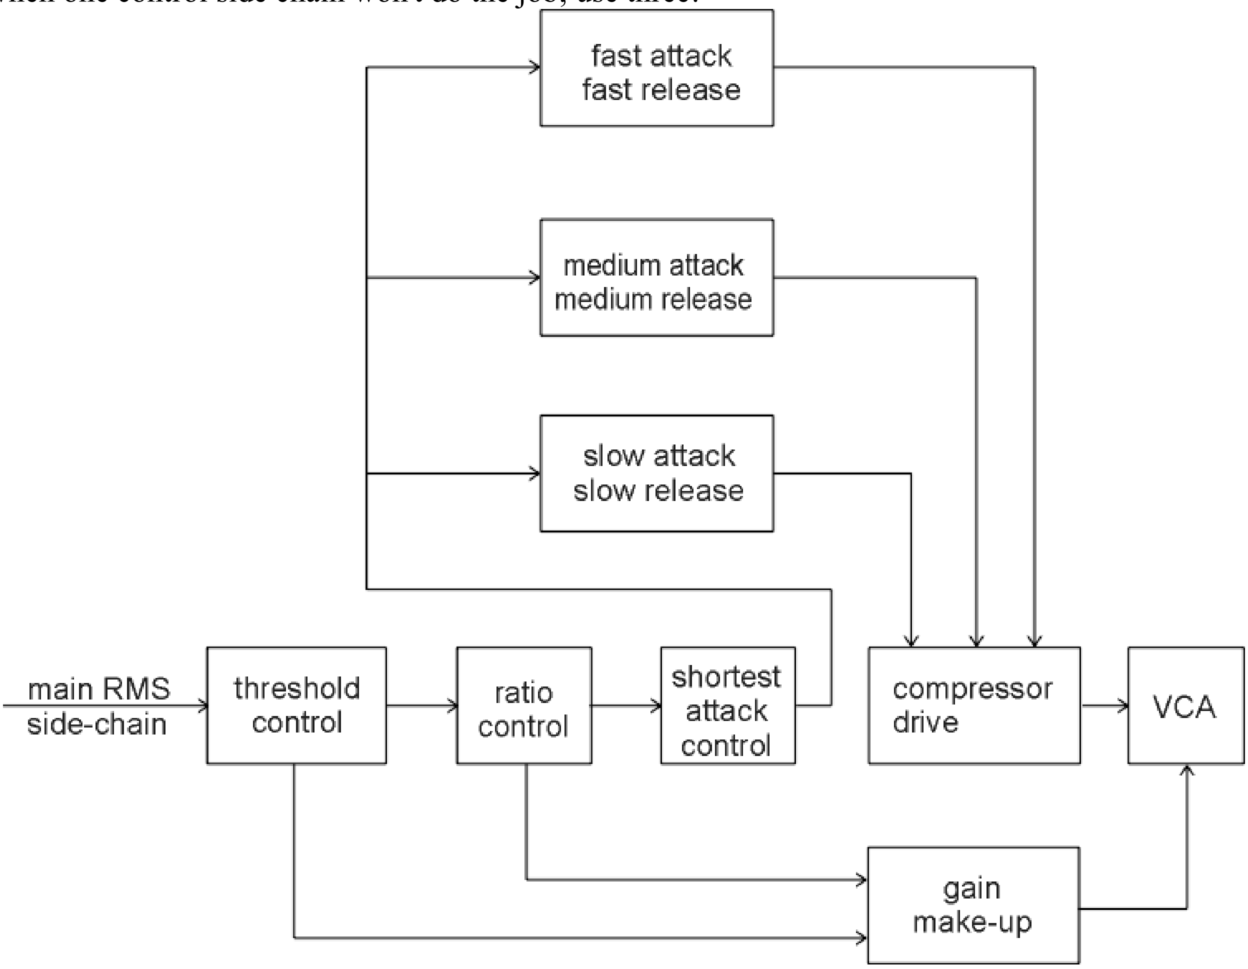

Side-chain 1 : a fast attack, fast release compressor which catches these over excited vocal phrases but recovers very quickly.

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 (slowest release time is only 500ms).

The Limiter : featuring multi-stage dynamic threshold control

We found many of the analogue limiters on the market fairly disappointing. Those which were fast enough to truly be called limiters sounded pretty awful when pushed hard into overload and those which sounded good were pretty hopeless at limiting. This is not surprising. 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 diagram below;

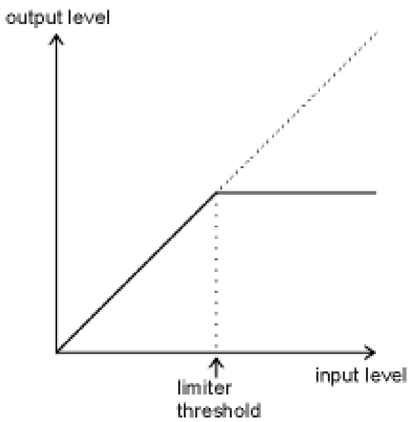

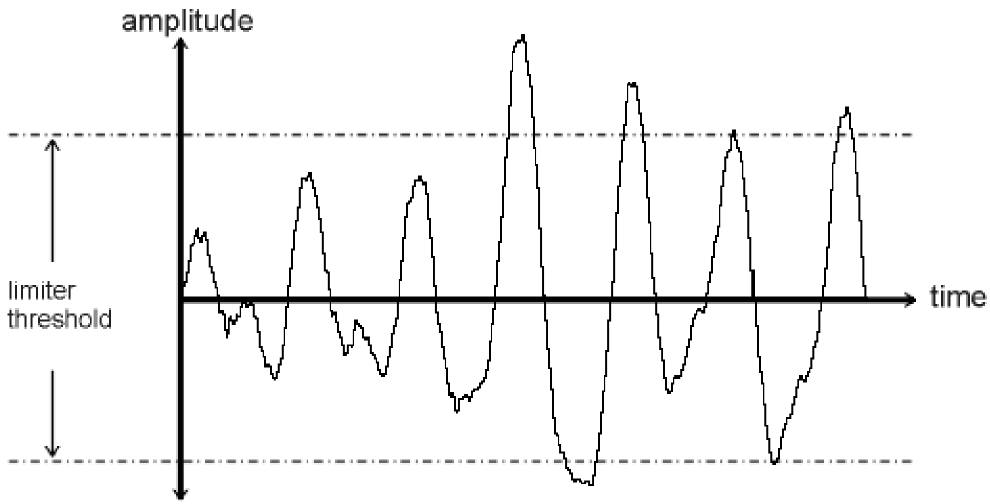

Audio is a time domain signal, measured as amplitude against time. If we consider a nominal limiter threshold; in the absence of an audio delay in the main signal chain, there is literally no time interval between an audio signal crossing the limiter threshold and, by definition, going into overload. 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. It won't do!

We had to look for another solution ....

In a live recording situation, a perfect limiter, if such a thing could ever be invented, would have to respond in advance of the problem existing and work out the best gain reduction characteristic for each and every audio overload. Even working in the digital domain it's almost impossible to achieve. In the analogue domain, working in real time, it really is impossible .... but we've tried!

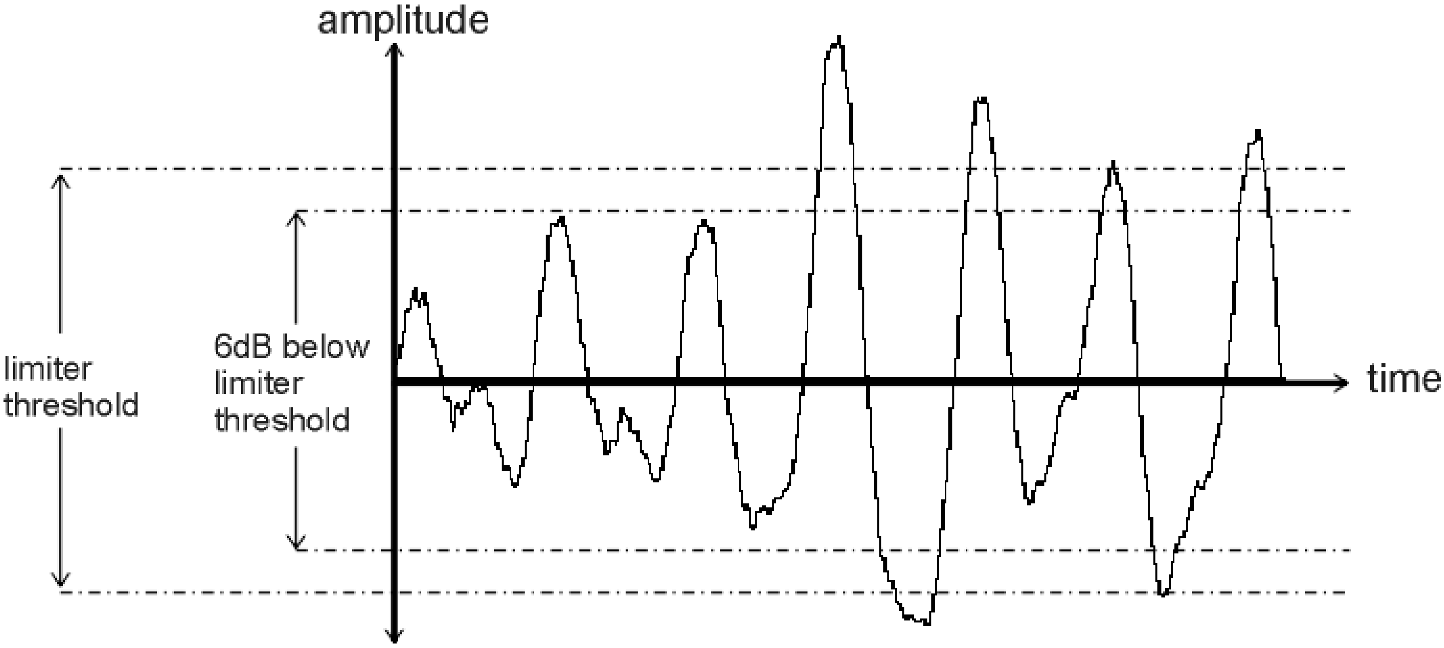

Although there is no way to guarantee the occurrence of an audio overload before it happens, it is possible to predict when one if likely to occur.

We decided to track the audio level not just at the limiter threshold but also when it crossed a threshold 6dB below the limiter threshold. Not only did we track the audio level when it crossed this lower threshold but we also measured the *acceleration* of the audio level increase. So fast rising edges score high on the 'accelerometer' and slowly rising edges score low. This is an excellent method of choosing what kind of limiter response curve is going to be best to tackle the potential overload of any audio sample. In some instances it also gives the possibility to tackle the overload just before it actually occurs.

So now we had a good tool to guide the limiter characteristic dynamically and in most cases we gain a little time to react just before the overload occurs.

Next we looked at the two most common problems of traditional fast limiters;

- Audio edges on 'live' material can be really fast. Even with a very fast traditional limiter (attack times in the region of 0.1ms, equals 100us or about 5 samples in a 48kHz sample rate recording system) we were able to easily overload a digital recording system and the overload was still audible.

  We raised the stakes, pushing the attack time down to 40us (2 samples), but we could still hear the overload on 'raw' live material, and in addition, the side effects of such fast limiting became all too apparent. Clicks, due to the very sudden gain changes, and very severe short term 'ripple' distortion.

  Through listening tests with various audio sources, we found that pushing limiter attack times much below 5ms began to audibly degrade the sound quality. Now, if the alternative is digital overload in the recorder then a fast limiter is preferable, but still not desirable.

- Audio distortion due to low frequency 'ripple' in the control chain. Most audio contains a wide variety of frequencies. Have a look at a typical audio sample (piano with sustain);

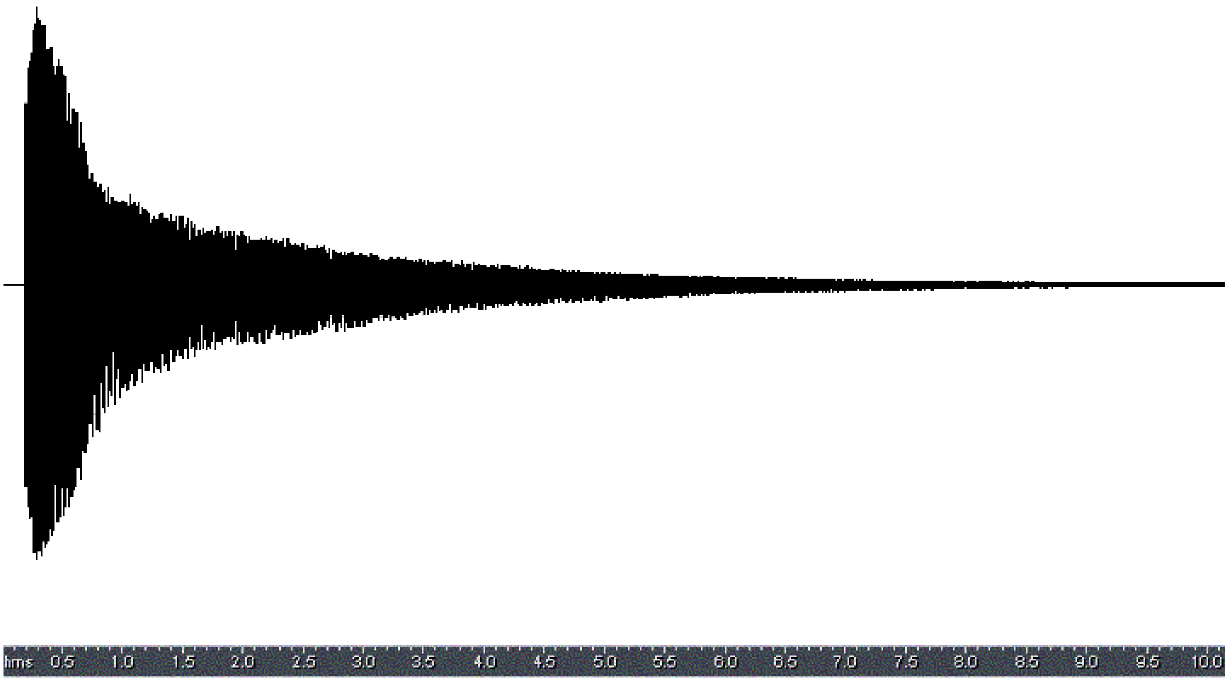

The sample has a total duration of around 10 seconds and the typical piano envelope can be seen, fast attack followed by primary and secondary decay cycles. If we zoom in to the first 500ms of the audio sample (below) it can be observed that the fast attack part of the envelope is actually made up from a number of more complex amplitude envelopes.

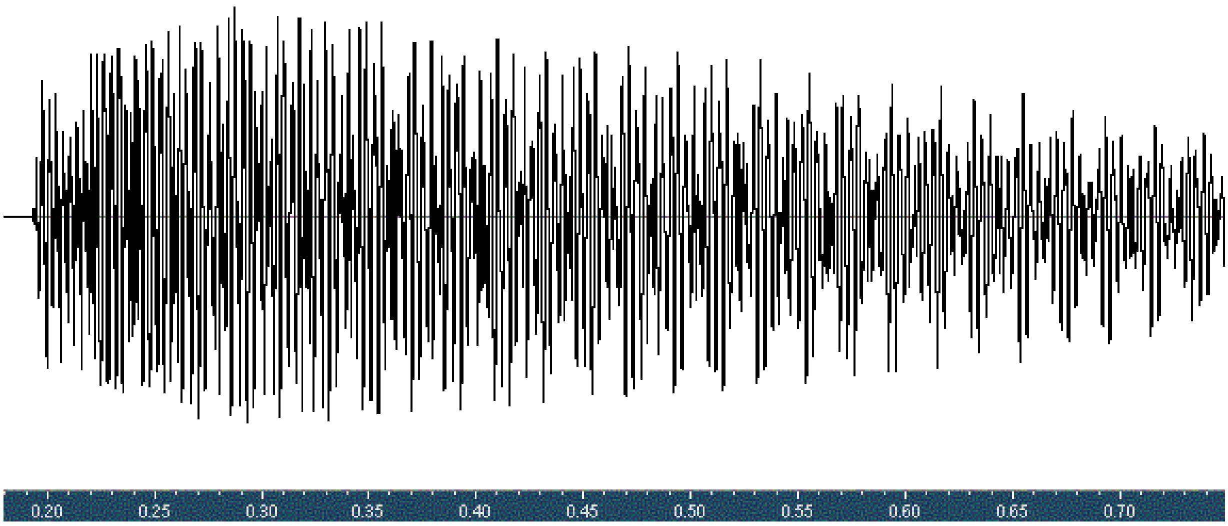

Zooming in further we can now see the individual sound components in more detail.

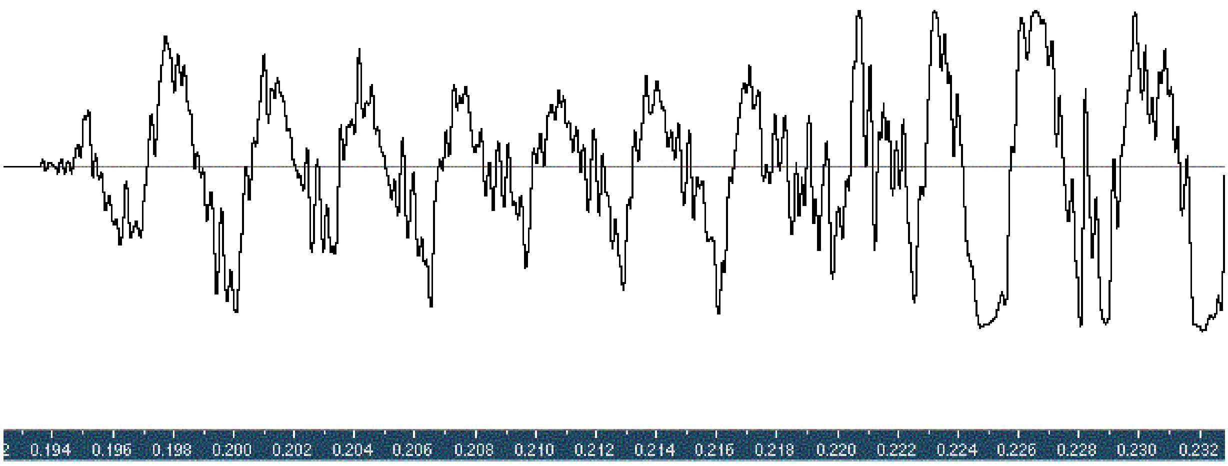

Most audio sounds of this type, whether they be instrument or vocal, consist of a fundamental pitch and a number of related harmonics. Additionally, there is usually a degree of tremolo (level modulation) and vibrato (small variations of pitch). In combination with the overall amplitude envelope, these give the sound its unique character.

Traditional limiter designs have a tendency to track this low frequency ripple which in turn modulates the level control element and causes distortion.

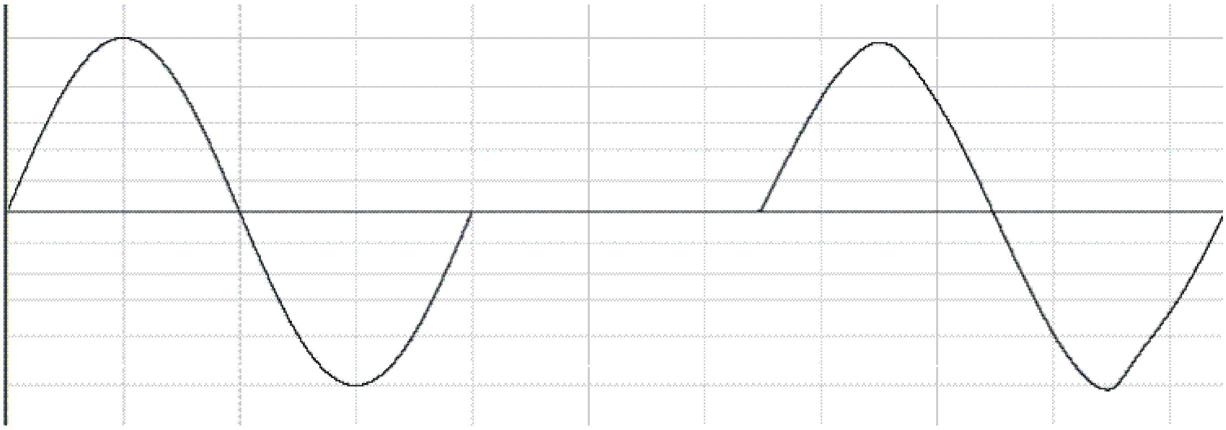

Sine wave test signal. Effect of low frequency ripple

The diagram above shows an example of this using a test signal. Notice that the sine wave is beginning to loose its smooth top and bottom curves and begins to resemble a triangular waveform.

Many audio sources including piano and vocals have either a primary low frequency 'carrier' component and vibrato/tremolo induced low frequency components. All can cause audible distortion in compressors and limiters.

One solution is to use an adjustable release time; longer release times tend to lessen this distortion but at the expense of adversely affecting the dynamics of the audio. Try a fast attack and slow release combination on a drum track and you'll hear what we mean. It just doesn't sound right! Whilst a manually adjustable release time in the limiter design would be a partial solution, life is too short! 'Sorry, can we do that take again. I need to adjust the limiter release control ....

So what did we do?

We started with a 5ms attack time limiter with a natural threshold of 0dB and then used the measured *acceleration* of the audio as it approached overload (between -6dB and 0dB) to dynamically reduce the limiter threshold by up to 3dB.

Simply put, the faster the audio was rushing towards overload, the more quickly we reduced the limiter threshold.

This has some key advantages;

- It has all the audio quality advantages of a slowish attack limiter but achieves the same limiting ability of a faster limiter. In this stage of our design we achieve the limiting effect of ‹1ms attack time but retain ALL THE SONIC ADVATAGES of a 5ms attack time. This has a significant impact on audio quality when limiting.
- The dynamic range of the recorded audio is optimised at all times so that audio which is 'gently' peaking to 0dB suffers little or no 'pre-emptive' limiting.
- As the dynamic threshold control is 100% analogue in nature, there is never a sudden transition into limiting; just like soft knee limiting, but totally responsive to the dynamics of the audio signal.

We were very tempted to apply this same process to a single sidechain limiter and increase the 'natural' attack time to around 0.1ms. The results were good, overloads were almost entirely eliminated and audio quality was very respectable; but not good enough; so we pressed on ....

Whilst trying to figure out a way to get that really fast attack time without audible distortion, we turned our attention to the low frequency 'ripple' distortion discussed earlier.

The limiter stage we had invented had a release time of around 100ms. This is set so that recovery from short overloads is fast enough to preserve the dynamic shape of the audio signal. It does however have a poor performance in longer periods of overload when ripple induced distortion occurs.

Simply increasing the fixed release time was unacceptable; this wrecked the original dynamics of some audio material especially, as expected, percussion but it also had a negative impact on some vocal phrases. During the limiter 'recovery' following a short overload the 'sparkle' in vocals was lost.

We decide to add a second side-chain to the limiter which would have all the benefits of our dynamic threshold control 'dtc' but would be engineered specifically to deal with ripple distortion. This second side-chain also has a natural attack time of 5ms; a longer release time of around 500ms; but with one significant difference. Whilst the 'dtc' process is applied progressively to side-chain one as soon as the -6dB threshold is passed; in the case of side-chain two, 'dtc' is measured from -6dB but only applied after a short time delay of around 25ms.

This delay in side-chain two has two advantages;

- It prevents side-chain two from contributing much during short duration overloads where it would be unnecessary, and undesirable to lengthen the release time (ripple distortion is a relatively slow time effect).
- It prevents side-chain two from overlimiting.

The combination of the two side-chains works really well and all but eliminates ripple based distortion.

But what to do about those really fast audio edges?

So far we had designed a medium speed limiter with adaptive threshold control providing very low distortion for both short and longer duration overloads and a very nice adaptive soft knee characteristic.

We decided that our starting point for the (hopefully!) final limiter stage had to be based on the dreaded audio clipper (fast but usually deadly to audio quality). The so called soft clipper showed the best promise. This is a technique which acts almost instantaneously on the audio waveform and progressively level reduces any audio which manages to escape the previous gain based limiter stages. It can react very quickly, but has one major disadvantage;

You cannot *gently round* an audio overload immediately it passes the soft clipper threshold; it needs some dynamic space to work in. Neither can you simply lower the threshold where it starts to work otherwise it soft clips below the threshold of the main gain based limiter.

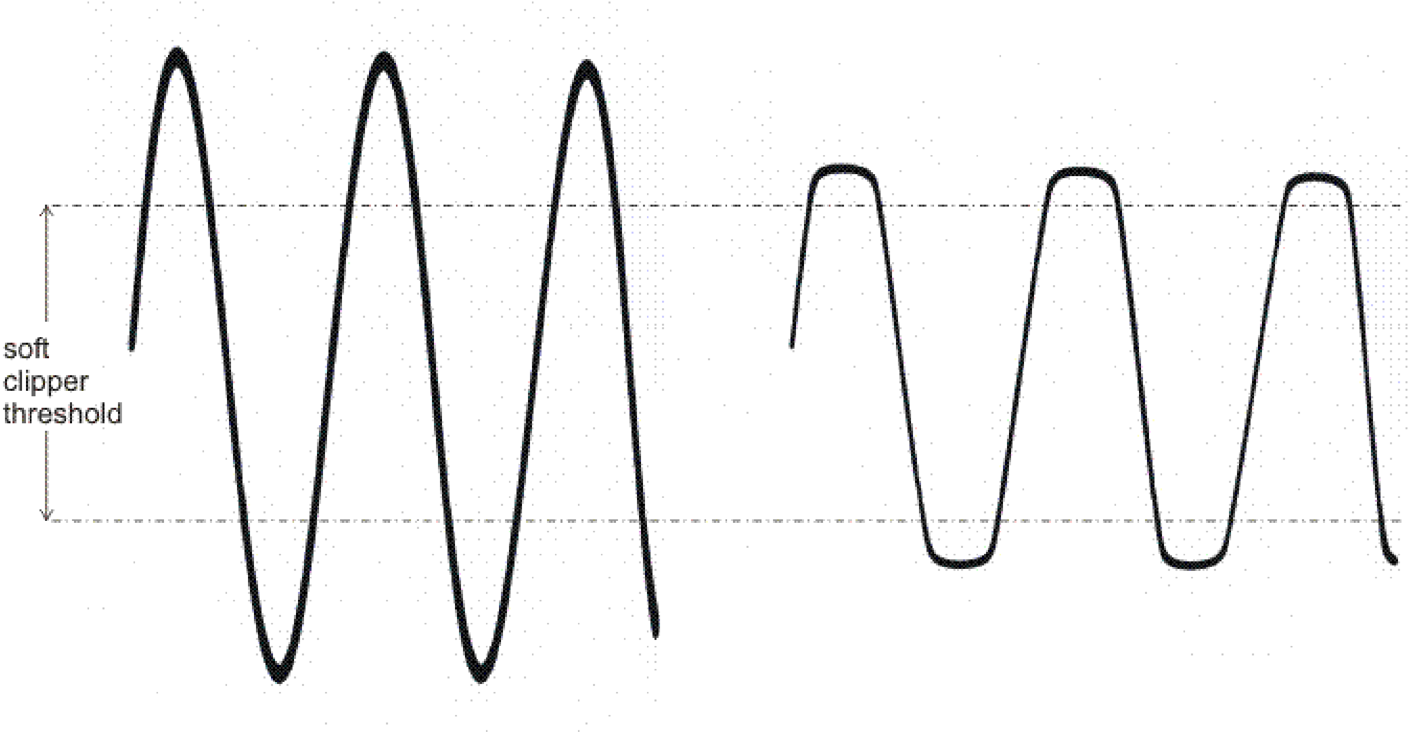

Test tone Result of soft clipping

Unfortunately, traditional soft clippers have to work within a very restricted dynamic range, typically 3dB above the normal limiter threshold, otherwise they would give no real protection against overloads which get past the main limiter stage. This small working dynamic range provides very little scope for 'smoothing' the audio peak and so audible distortion can be a problem with large audio overloads into the soft clipper stage.

This is still preferable to overloading a digital recording system but badly degrades the audio quality for anything but very short overloads.

So we designed a final soft clipper stage which deals with any short duration overloads which escape the previous sidechains one and two and we gave it a more generous 6dB window to operate in (this greatly improves the distortion performance). But hang on! haven't you just stolen 6dB from the recording device's dynamic range?

Then we applied dynamics threshold control to this stage as well!

So now the soft clipper can operate dynamically adjusting its threshold to suit the overload which has dared escape the previous two limiter stages. It can limit fast edges to the maximum desired level, but it has enough dynamic range to work within to be gentle.

This final limiting stage typically operates on only one or two consecutive overload audio peaks. This is enough time for side-chain one to get up to speed and take over. Side-chain two gets into action only for longer duration overloads to protect against ripple based distortion.

The complete design scheme is shown below;

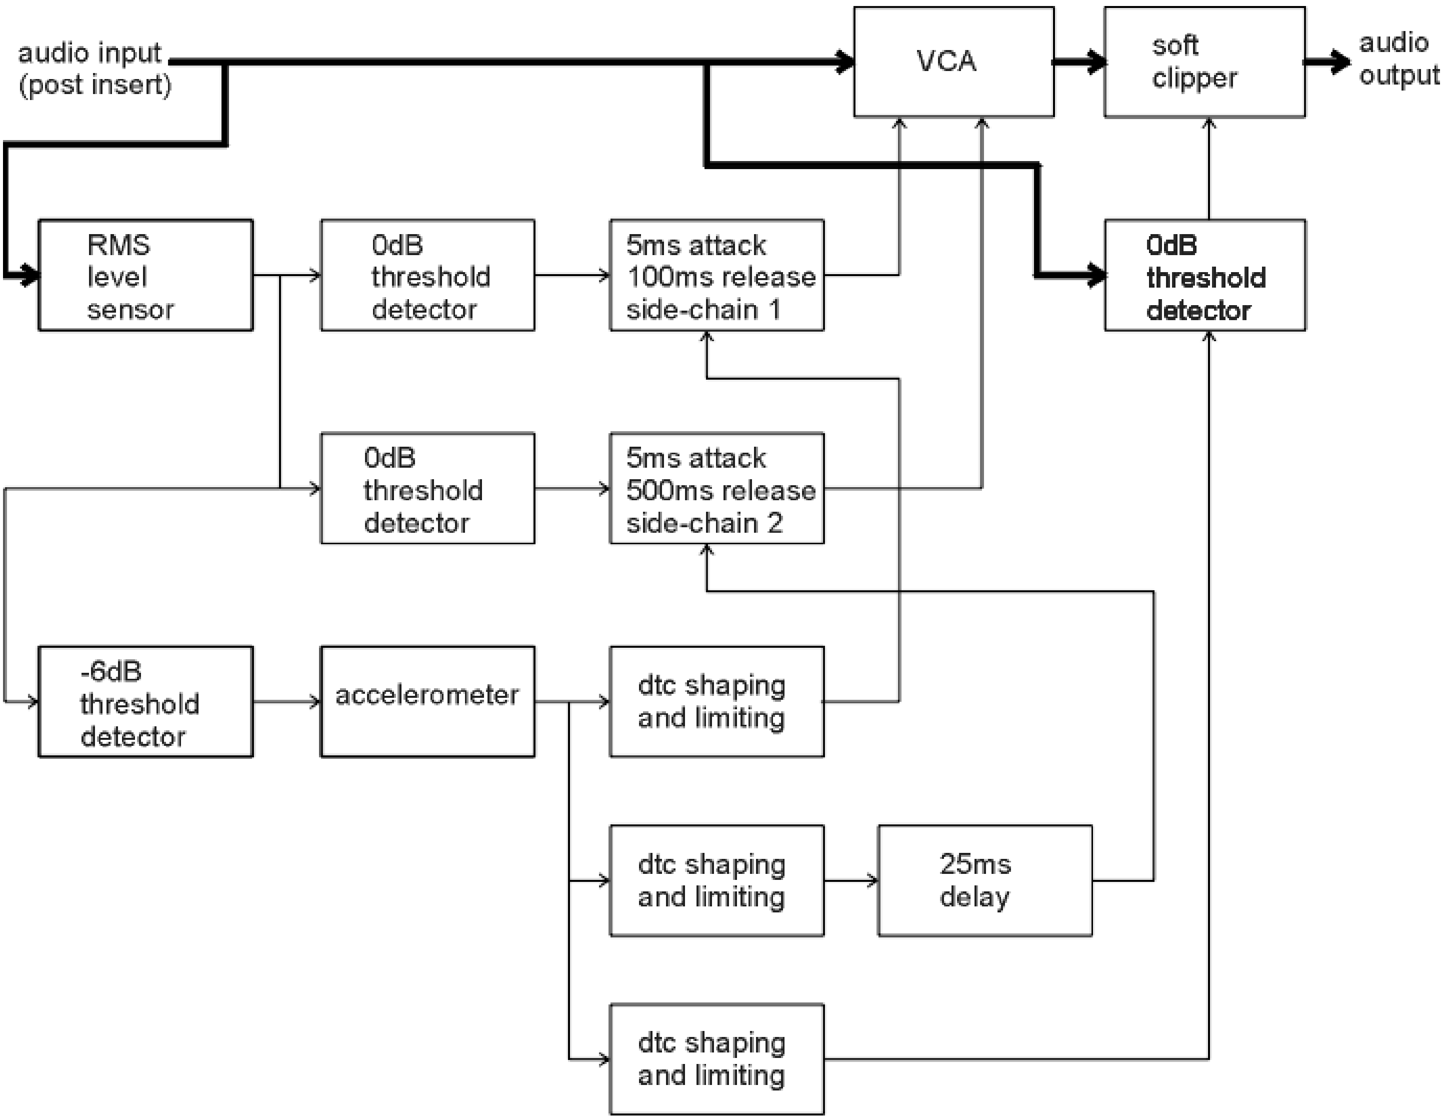

Outline design of dynamic threshold control limiter

Resulting performance of the P1 limiter

This limiter design as implemented in the P1 processor will limit a very fast rising audio edge to within 0.5dB of the limiter threshold with an audio input +6dB above the limiter threshold.

Serious overloading of the limiter can push the overshoot beyond 0.5dB but with sensible setting of the P1's gain control and setting up the digital recording system as described in the user handbook, it should be possible to record without digital overload with most types of source material.

It would have been possible to add a further hard clipper stage to prevent any audio overload clipping the digital recording system but we think it is preferable for the user to be able to see any hard clipping as 'overs' in their recording system.

Perfect? .... Of course not! but we've gone a long way to eliminating the undesirable side effects of many traditional analogue limiter designs.

Could this all be done in the digital domain? Probably; but then it wouldn't be analogue .... would it?

© Safe Sound Audio 2003