ANALYSIS AND SIMULATION OF ANALOGUE DYNAMIC COMPRESSORS AND LIMITERS IN THE DIGITAL DOMAIN

Michael J Kemp

Sintefex Audio Lda
S. Marcos da Serra, P-8375 Portugal.
+351 282 361748 (fax +351 282 361749)
mikekemp@sintefex.com, www.sintefex.com

PRESENTED AT THE 109th AES CONVENTION, LOS ANGELES SEPTEMBER 22-25, 2000. (preprint #5185)

Abstract - A method is described of applying non-linear synthesis to analyse and simulate traditional audio dynamics compressors and limiters. A set of level dependent impulse response measurements is taken at each of multiple attenuation levels. In addition a measurement of attenuation characteristic against applied signal amplitude is made at various settings of the device and applied in simulation and a bilinear dynamic convolution is performed varying on a sample by sample basis.

I. INTRODUCTION

The audio recording engineer is repeatedly faced with the requirement to reduce the dynamic range of material being recorded and this has been achieved through the years with various analogue approaches, for example variable mu vacuum-tube devices, photo-electric light dependent resistors, balanced semiconductor gates, field effect transistors and even mechanically driven potentiometers. All these analogue techniques have various problems associated with them, for example problems with distortion, noise and linearity of gain variation versus applied signal.

With the advent of digital audio processing one would initially have thought that all these problems could be consigned to history. Digitally it is possible to apply any conceivable compression algorithm to an audio signal, with any desired degree of (or lack of) distortion and noise, and any desired attack and decay characteristics, limited only by the delay permissible through the device. Indeed in off-line processing this is no problem and an audio recording can be analysed and its gain precisely adjusted sample by sample to any desired characteristic.

However, in practice digital compression is usually less than satisfying to the ear. Two reasons may be postulated for this.

In the first instance original audio material can exceed in dynamic range that of the initial analogue to digital converter in a system, requiring dynamics compression to be applied in the analogue domain. Further digital compression is likely to raise the low level artefacts of conversion. However we are now in the position that original conversion can approach 120dB of dynamic range and be free of artefacts by properly designed dithering algorithms and it is hard to imagine an original signal source that would exceed this in dynamic range.

A second and deeper problem with digital compression is related to complexity. Digital compression is usually felt to sound flat and uninteresting, despite the phenomenal control we have over the compression algorithm. This may be acceptable if the requirement of the compression is just to meet some requirement of the transmission channel, for example to prevent overload of a transmitter, but when used as a creative medium to interpret an original audio source into a recording we frequently find ourselves returning to analogue and dusting off old analogue compressors to achieve a familiar and desirable sound characteristic.

It should be mentioned that even in the case of pre-transmission signal processing, a characterful compression is desirable to optimise signal level while retaining a musical quality to the signal. This is often not achieved with digital multi-band compression which although offering high average signal levels can also lead to audience fatigue and switch-off. Here the station managers are making a creative decision relating to the quality of their station output, and to keep their audience they must consider the aesthetic effects of the processing.

There are a number of practical problems with returning to analogue during increasingly digital recording sessions. Analogue equipment, particularly older items, are quirky; the performance is apt to change with time and temperature, settings cannot be stored and recalled, multi-channel units cannot be matched and are frequently not available in the configurations needed for multi-channel recording formats. An effective digital simulation would address many of these problems and allow a gallery of effects to be available for quick auditioning.

Although there have been attempts to model analogue processors this paper presents a method whereby it is possible to capture the key performance characteristics of existing analogue dynamics controllers and to save the information for repeatable reuse wholly in the digital domain.

II. COMPLEXITY

A significant feature of an analogue processor is that the effect is complex. Analogue circuitry passes the signal through a gain control element, the performance of which varies with the amount of gain reduction, and there are various amplification stages and often transformer coupled inputs or outputs, each of which adds to the response of the device. Analogue design has tended to choose the approaches that produce acoustically pleasing results, and of course the successful ones have survived and are now much sought after.

It is not straightforward to simply code in complexity when designing a digital compressor. However it is practical to consider sampling some of the characteristics of analogue processes and to reproduce these in their natural complexity in a digital simulation. Three main characteristics of an analogue dynamics compressor are (i) the main signal path performance (at one or more gain reduction settings), (ii) the compression curve of the device, and (iii) the dynamic performance of the device. The approach taken here is to look at the first two of these areas and analyse and simulate them, and to provide a road-map for simulating the third.

Snapshot or Variable Simulation

A point worth mentioning here is that there are two main requirements of the analysis and simulation method. A “snapshot” simulation represents the capturing of a signal path with predefined settings. An example of this may be an analogue mastering signal path which achieves some desirable effect with a combination of equalisation and compression.

A variable simulation requires multiple samples of a device and provides the user of the simulation a set of familiar controls which can vary the effect in a similar way to the original device.

These requirements have many similarities but each has its own special issues. The discussion that follows focuses mainly on the variable simulation. One issue here is that it is possible to provide more options in the simulation than in the original device. Some people consider that this should not be permitted and that a simulation should be limited to the options available originally. However this is a philosophical issue that does not need to be further discussed here.

Static Audio Performance

The performance of an analogue dynamics compressor may be measured in a number of ways. For example a frequency response can be measured, as in fig.1, taken at unity gain (no gain reduction) of the device under test (DUT). As a comparison, the frequency response at 10dB of gain reduction shows a marked change (fig. 2).

A measurement of total harmonic distortion plus noise (THD+N) at unity gain shows fairly significant distortion at around -55dB (greater than 0.1%), with an input signal of 17dBu.

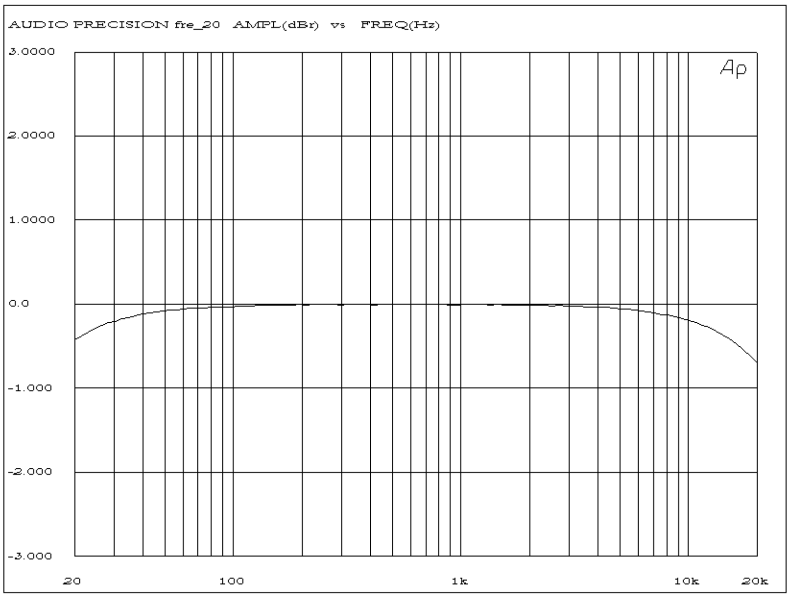

Fig 1. Device Under Test (DUT) unity gain frequency response

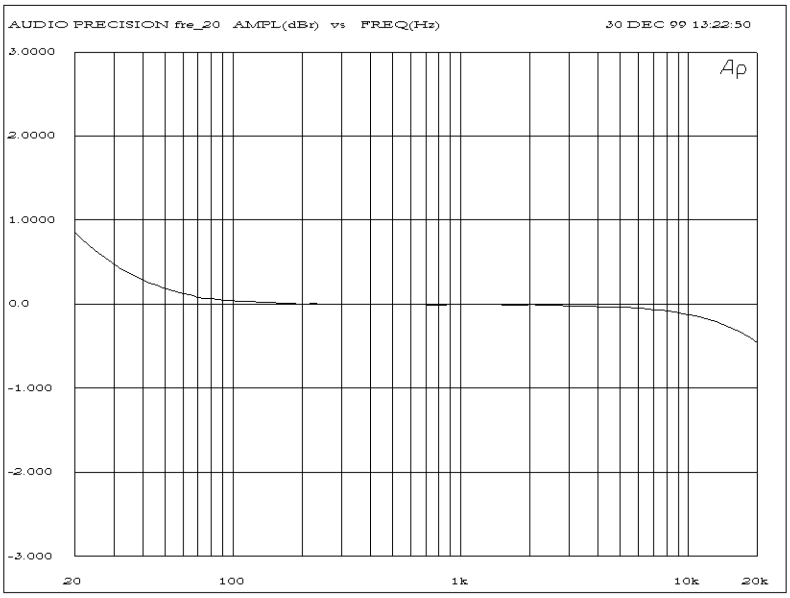

Fig. 2. DUT at -10dB gain.

These examples indicate that the device will have a characteristic sound, based both on the changing frequency response and the distortion characteristics. Audio engineers will be familiar with the sounds of a variety of compressors and be aware of material that suits each best.

Dynamic Performance

In addition to the static performance of these devices there is also the dynamic effect which is dependent on (a) the time constants of attack and decay, and (b) the compression curves of the device. These are covered in more detail below.

III. DYNAMIC CONVOLUTION

The author has already described [1] a process for capturing the characteristics of an analogue processing chain of fixed gain and this process is now referred to as Dynamic Convolution.

Dynamic Convolution involves first taking multiple impulse response measurements of an analogue processor at a sequence of different levels. A new signal may then be processed using this set of responses in a modified convolution algorithm which dynamically selects the appropriate response to apply dependent on the instantaneous sample level of the new signal. This selection is done on a sample by sample basis and in addition interpolation is performed between impulse responses representing adjacent impulse levels on either side of the new signal sample level.

Dynamic Convolution, when applied to a new signal, results in (i) the signal being subject to the same impulse response as the original equipment, implying the same frequency and phase response, and (ii) the variations due to non-linearity in the original equipment being simulated, resulting in signal dependent variations in harmonic distortion. Listening tests indicate that the characteristics of the original device are captured well and frequency response measurement show this to be born out in practice. The accuracy of simulation depends on (i) the degree by which the original signal path deviated from a linear response, (ii) the number of different amplitude impulse responses taken, and (iii) the number of steps in the convolution.

For the purposes of this paper, the dynamic convolution algorithms were performed in real time using the proprietary hardware design embodied in the Sintefex Audio Digital Audio Effects Replicator (fig. 3). A block diagram of the principle processing elements of Replicator is shown in Fig. 4. This unit comprises multi-channel processing ability with each channel provided with 4 cascaded 60MHz floating point digital signal processors (DSPs) as the dynamic convolution pipeline, allowing for a 2048 step interpolated dynamic convolution at 48kHz on each channel. Each pair of channel processing pipelines is controlled by a further DSP handling input and output, input sample evaluation and real-time selection of impulse response pairs and interpolation factors, as well as providing further signal processing options. The system of up to 8 channels is further coordinated by another DSP accepting user input, distributing impulse response data from the disk storage and deriving control values.

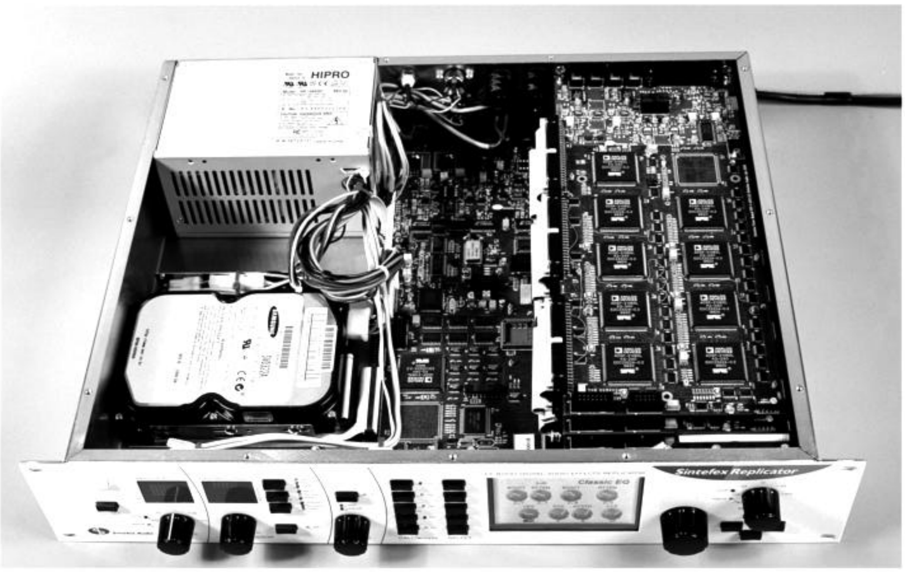

Fig 3. Simulation hardware showing a dual channel processor chain from an 8-channel unit.

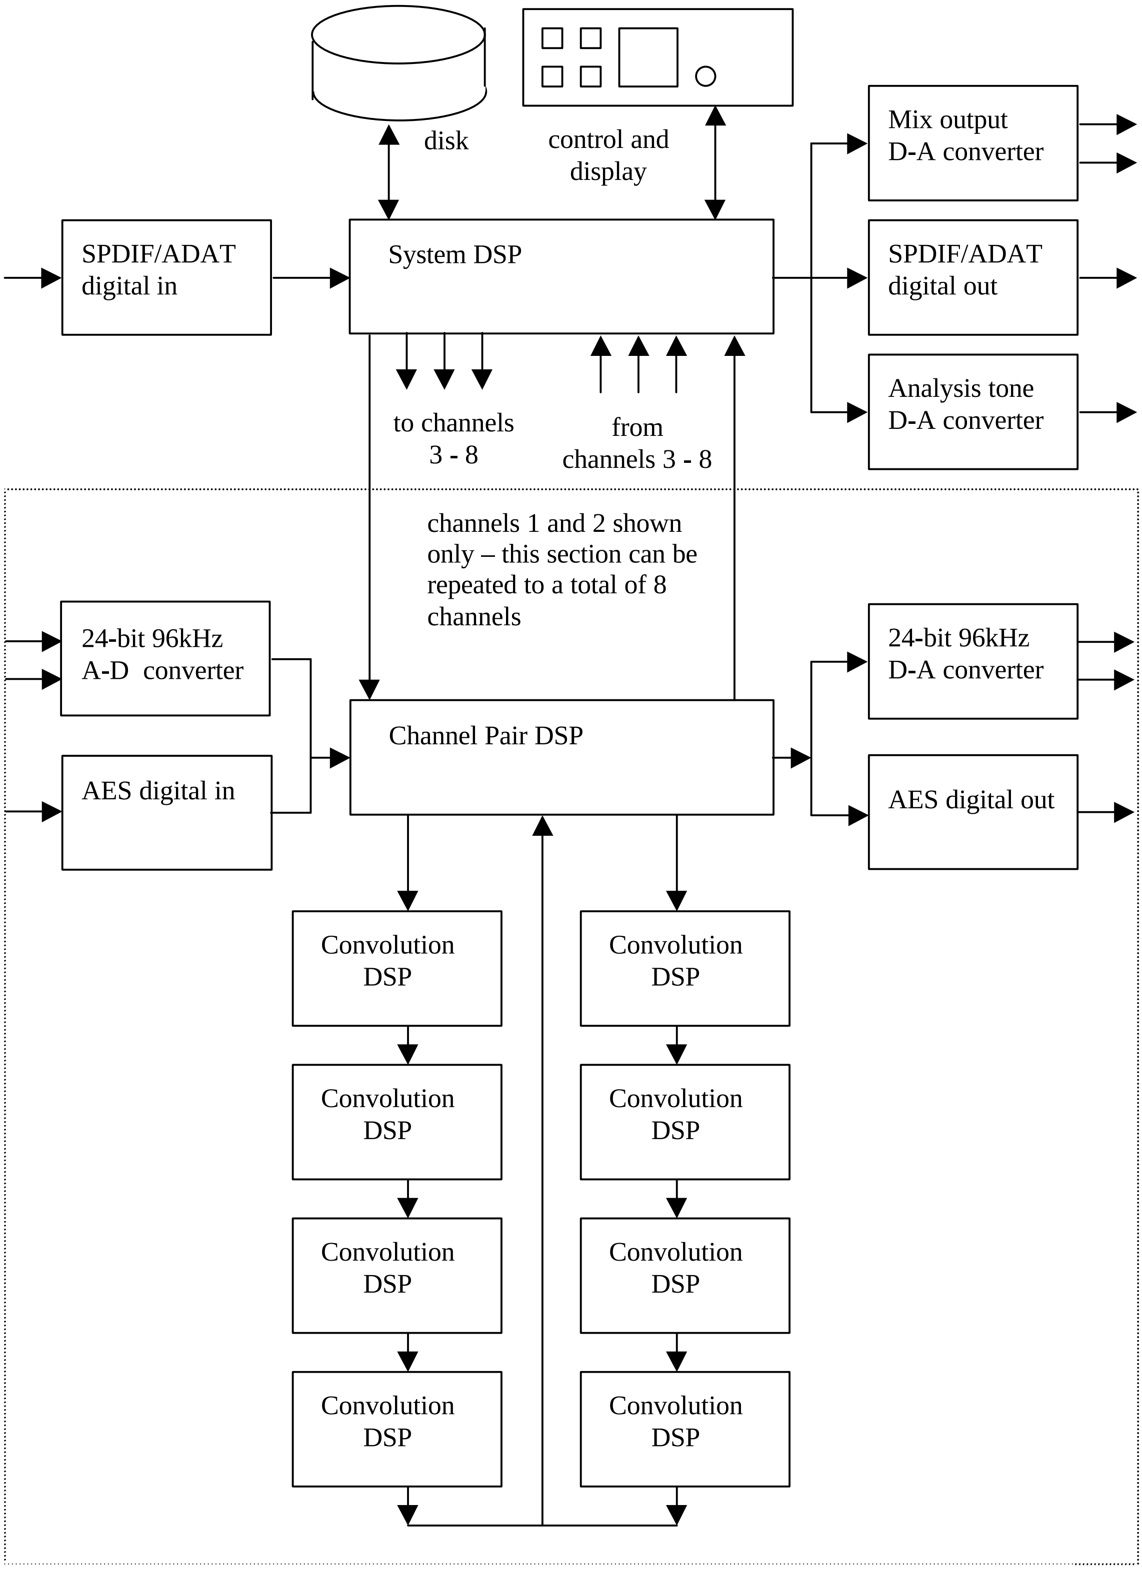

Fig.. 4. Block Diagram of Principle Processing Elements

Fig. 5 shows the frequency response of the simulation of the compressor path at unity gain (as shown in fig. 1) superimposed on the original frequency response. This uses a 2048 sample convolution at 48kHz and a linear ramp window on the impulse response.

Fig. 6 shows an illustration of the impulse response data for the device taken from the Replicator display. It does not show the low level signal detail and has been expanded so only a few tens of samples are on the screen. The groups of 8 responses shown are a subset of the 128 different level impulses taken from the device under test. Although this does not reveal detail it can be seen that there is some complexity in the impulse response both before and after peak of the response to the impulse stimulation.

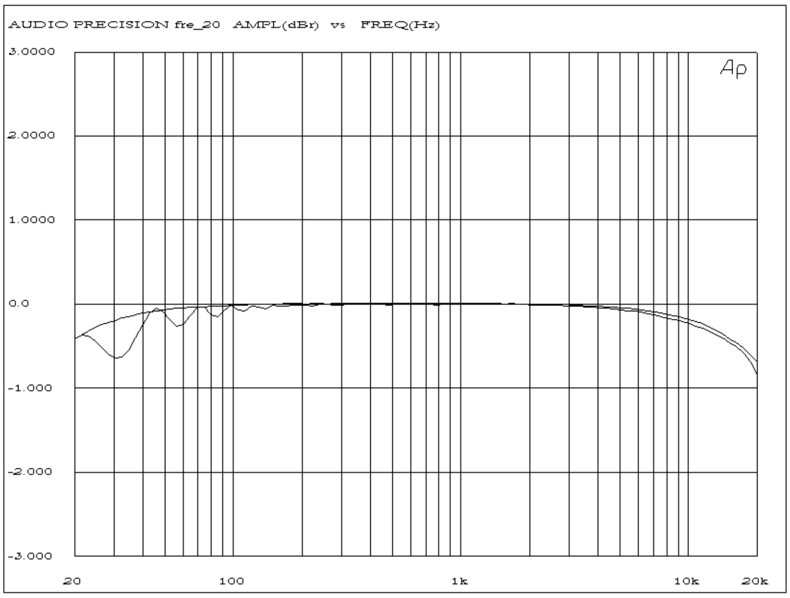

Fig 5. Simulation of DUT unity gain response (2048 samples, 48kHz), superimposed on original DUT response.

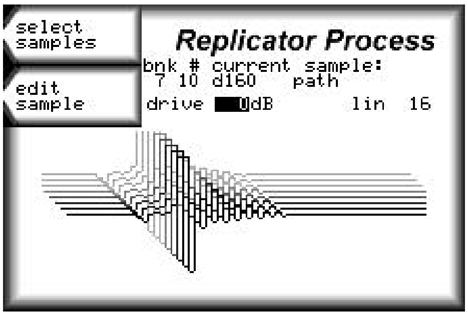

Fig 6. Example impulse display, DUT simulation, unity gain

The simulation process of dynamic convolution may optionally be switched off. In this case the convolution becomes linear and although frequency and phase response is simulated there is no simulation of the distortion characteristics of the original equipment. Measurement of the simulation of the DUT of figure 1 shows THD + noise at around -90dB in the linear mode (limited mainly by the converter noise) and -60dB when in dynamic convolution mode, modelling the device under test.

IV. COMPRESSION CURVES

The second area of complexity that analogue processors provide is in the characteristic curve of the compressor, that is the output level expected from any given steady state input level. In creating a digital compressor we can define any curve we desire, but when we choose to examine the curves of well known analogue compressors we find a interesting variety of at first sight non-ideal curves. Once again these lead to distinctive and desirable properties of the analogue device. The problem to be addressed is how best to store this characteristic, and it is best to look at the method of simulation before considering this.

Simulating Curves

After some consideration it was decided that rather than storing compression curve characteristics algorithmically, it was more flexible to adopt a table based solution. A table of gain reductions is required and it was decided to use 1dB input steps, typically from peak signal level down to at least 85dB below this. A set of such tables is desired representing a number of compressor slopes.

It is then possible to simulate each of these curves by deriving a further linear gain table that may be directly addressed in real time. This can be interpolated from the gain table in decibels and has the benefit that there is a lot of accuracy at the high end of such a table (half is devoted to the top 6dB of input level) and the “granularity” of the table is still at the 1dB step level down to the minus eighties of dBs.

Sampling Curves

Sampling curves can be carried out in a straightforward way by applying a sine tone at a desired “probe” level and monitoring the output from the device under test. The level is allowed to stabilise and then a direct reading of the gain through the device under test can taken. The process can then be repeated at 1dB steps of the probe signal until a direct list of gain at each input level is achieved.

The gain is clearly subject to all the manual gain control elements of the device under test as well as its threshold setting. The manual gain settings can be effectively removed by defining the gain applied to the minimum level probe signal as “unity” gain from the gain control element, and normalising the table by subtracting this value from all values in the table. In simulation it is possible to add arbitrary gain reductions as desired, or the original figure may be restored if is desired to simulate a particular “snapshot” of the device under test at the exact settings in effect at the time of the test.

If it is desired to simulate the whole range of the device under test, a whole family of curves can be derived at different compression ratios. It is not normally necessary to sample at different thresholds as this can usually be derived by considering threshold as a combination of adjusting input gain and output gain.

The threshold of a particular sample is recorded so that simulation can take place in a consistent ways between simulation of different machines. The sampled gain data is stored in the master gain table so that the threshold is at a known level. The sampled threshold may (i) be assumed from the control labelling on the original device under test, (ii) be inferred from auditioning and visual inspection of the curve data, or (iii) be derived by programmed analysis of the data values. We have so far relied on (i) and (ii) but clearly the method of (iii) can be implemented in due course.

Interpolating Curves

When a set of compression curves has been taken from a device under test, it is possible to offer not just those curves but to generate interpolations between them. Indeed, if interpolated with an ideal 1:1 curve, a continuous range of compression ratios down to 1:1 can be offered.

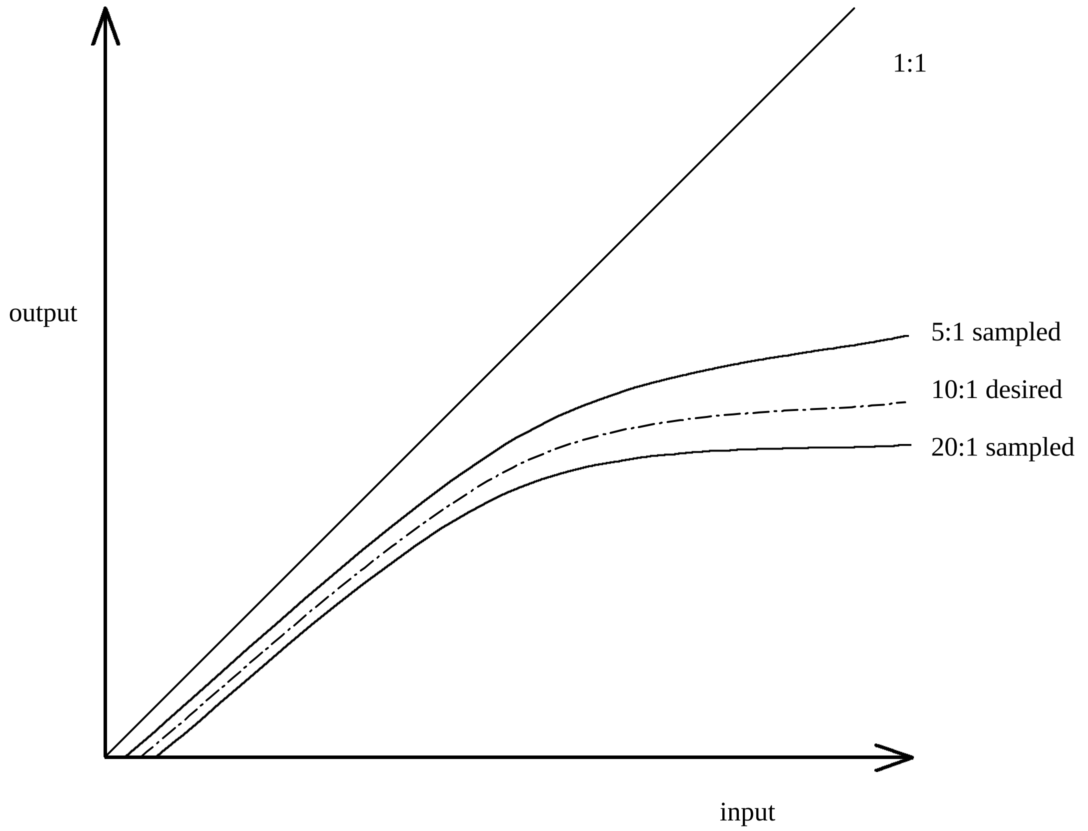

Fig 7: Interpolating 5:1 and 20:1 curve to get a 10:1 curve

Fig. 7 shows an example where a 20:1 curve and a 5:1 curve has been sampled from a device under test. The dotted 10:1 curve can be interpolated by deriving a coefficient of interpolation α using the equation

$$\alpha = N(M-R) / R(M-N), \qquad (1)$$

where N is the lower ratio N:1, M is the higher ratio M:1, and R:1 is the desired ratio. In this example, M = 20, N = 5 and R = 10, so α =1/3.

It is then possible to derive the desired gain reduction at any given input level from the equation

$$Z = (1-\alpha)X + \alpha Y, \qquad (2)$$

where X is the gain reduction from the higher ratio curve, Y is the gain reduction from the lower ratio curve and Z is the desired gain reduction for the interpolated curve.

Interpolation towards a 1:1 curve may also be performed using these equations with N =1 and Y = 0.

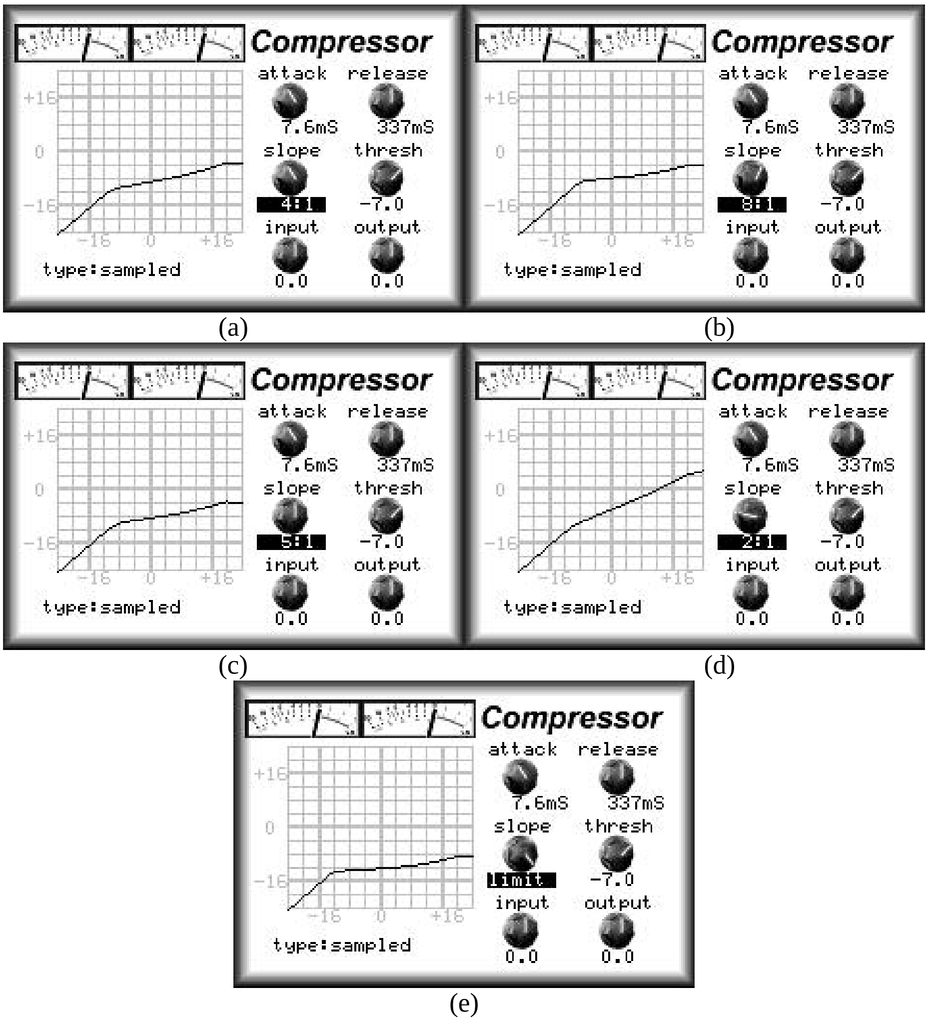

Fig. 10. Some sampled compressor curves taken from simulator display.
(a) 4:1 curve, (b) 8:1 curve, (c) 5:1 interpolated between 4:1 and 8:1, (d) 1.5:1 interpolated from 4:1
(e) "distressed" setting of compressor sampled with "all buttons down"

Example Curves

Fig. 8 shows some screen shots from the simulation of sampled curve from a well known compressor which exhibits some interesting gain characteristics. The curves show input level along the horizontal axis and corresponding output level along the vertical axis. Fig. 8(a) shows a sampled curve at 4:1. It can be seen that the knee of the compressor is fairly sharp and that as the signal rises above the threshold the slope of the curve "eases off" to a lesser slope, estimated at about 2:1. The second knee at the right represents the peak level sampled through the device under test, beyond which we choose to impose a straight line limit on the output level rather than extrapolating the curve from the sample data. Since this point is about 30dB above the knee this seems to be a reasonable approach.

Fig. 8(b) shows a sampled curve at 8:1, and it is interesting to observe that the threshold of the device has changed. Both samples were taken with all controls except the compression ratio at the same settings. This sort of control interaction is of course part of the nature of analogue devices.

Since we are concerned with preserving the quirks of these devices we do not try and adjust out these variations, but simply assign them a slope as defined by the device under test. The operator may select a slope not sampled from the device, as for example in fig. 8(c) which shows an interpolated 5:1 slope. In this case the curves for 4:1 and 8:1 are used in the calculation already describes to derive the 5:1 slope which therefore shows some spreading of the knee due to the threshold shift in the device under test.

Fig. 8(d) shows an interpolation between curve 8(a) and unity ratio to generate a 2:1 curve. The accuracy of all these curves is of course limited by the accuracy of the original device.

The compressor under test here featured pushbutton selection of compressor ratio. Fig 8(e) shows a curve derived from the compressor when all the ratios are selected simultaneously. It can be seen that the threshold drops yet again and the curve becomes effectively flat, or limiting. This setting therefore instantly increases gain reduction and makes the effect of the compression more dramatic and is in fact often used in extreme cases during music recording. Here we have assigned this as a “limit” curve which may be selected for simulation.

V. SAMPLING THE SIGNAL PATH

We have shown the results of taking an impulse response set of a typical device under test at unity gain and applying it in a simulation to a new signal (fig. 5). This may be simply achieved using the techniques of [1] provided that the device under test may be set to a compression ration of 1:1, or the threshold of the device under test can be set sufficiently high that no gain reduction starts to occur even during the high level impulse section of the impulse sampling procedure.

However a problem arises when the controls of the device do not permit this to be achieved or it is not possible to alter the controls of the device from a predetermined setting. This latter situation may occur when it is desired to sample a critically adjusted device where it might be considered impossible to reset the device to the desired setting after sampling.

The dynamic convolution process requires a set of impulse response samples of the device taken at different amplitudes, however the characteristics of the device must not change during this procedure. Since a compressor would tend to alter its gain it is necessary to circumvent this.

Since a mechanism has already been created to assess the gain curve of the device under test it is straightforward to apply a tone of a known amplitude to achieve a desired gain reduction from the device under test. This tone can then be removed and the sample impulse applied. It will be known that at the time of applying the impulse the gain of the device will be the same regardless of level of impulse, and as long as the time constants are not set to too rapid a release the gain will be very close to that desired. In this way a consistent set of impulse responses can be achieved to simulate the signal path at any desired gain reduction. It is also necessary to normalise each signal path response to provide the same overall gain for use in the simulation.

VI. DYNAMIC GAIN VARIATION AND BI-LINEAR CONVOLUTION

The compressor simulation has already been described by which the gain appropriate to any given level is derived in real time from a look-up table. After application of appropriate attack and release time constants this may be used directly as a multiplier to apply the gain reduction digitally.

In addition it is possible store a number of signal path characteristics representing different gain reductions of the device under test. The calculated gain is used to select two signal path characteristics representing gain reductions on either side of the gain to be applied at any instant, together with an interpolation coefficient for linearly combining them in proportion to the position of the applied gain between these two sampled signal paths.

The dynamic convolution is then used to apply these two convolutions in parallel, combining the result. This represents a convolution with two orthogonal linear interpolations being applied (hence bi-linear), one dependent on the instantaneous amplitude as required by dynamic convolution, and the other derived from the instantaneous gain modification required.

In situations where the gain required falls outside of any sampled gain characteristic, the nearest is used.

In practice a simulation with a single sample of the gain path of the compressor is sufficient to characterise some devices. However where there is significant change of response as gain reduction changes it becomes necessary to store two or three signal path characteristics at, for example, unity gain, 6dB of attenuation and 20dB of attenuation, to simulate the way the sound of the compressor device changes as the gain reduction changes.

VII. FURTHER WORK

To complete the simulation of signal paths with gain compression built in, especially where it is desired to "snapshot" a desired analogue signal path, there are few more characteristics to be assessed and included in the simulation. These have not been incorporated at the time of writing this paper but the mechanisms have been defined and are summarised here.

Assessing Time Constants

Attack and release time constants can be directly determined by use of the sine wave probe signal used to map the compressor curve of the device. Since the gain of the device is being determined by analysing the return signal it is also possible to estimate the time constant when the probe signal is changed to a different value. This can be stored and simulated using the time constants of the digital simulation. Further analysis can also reveal any dual time constants employed in the original equipment.

Assessing Side Chain and Precompression Equalisation.

In a pre-configured signal path it is quite likely that there will be equalisation before the gain reduction processor, or equalisation in the side chain of the compressor. If all the equalisation is known to take place before the compression the simulation can be done by placing the dynamic convolution process which simulates the entire equalisation of the signal chain before the gain reduction determination element of the simulation. Similarly if all the equalisation takes place after the gain control element this is easily simulated.

Where there is some equalisation both before and after the gain controlling element, or if there is side-chain equalisation, it is required to make an estimate of the combined effects of these two elements and insert a matching equalisation into the side-chain of the simulation. If it is possible to access a side-chain output from the compressor this can be measured directly, either by taking an impulse response of the path, or by applying a sine wave analysis signal at spot frequencies across the audio band. An impulse response can be used for direct simulation using a convolution but it is computationally more efficient to approximate the equalisation from signal input to the gain control element by the use of infinite impulse response (IIR) filters calculated to give a reasonable approximation. Since this signal is used solely for deriving gain reduction it is not necessary to simulate every nuance that is required for the direct signal path.

Where the side-chain signal cannot be accessed, the equalisation of this signal path can be inferred by measuring the actual gain reduction curve as already described, and repeating this procedure for a set of spot frequencies. An IIR filter can then be interposed as above in the gain determining calculation to simulate this effect.

In any of the above cases, the whole equalisation path is sampled as has been described to achieve an accurate simulation of the overall signal path, as varied by the amount of gain reduction.

Expanders and Gates

Clearly the techniques discussed here can be easily extended to the analysis of expander devices. Gate devices will require a probe signal that both decreases and increases, so that the gain curve can be inferred by reference to the output signal level and the "hysteresis" effect of the gate can be estimated. This can be done algorithmically or by duplicating the method used in analogue gates and expanders of swapping the gain sensing input from before the gain reduction element to after, and generating a lookup table of gains based on output levels.

VIII. SUMMARY

An effective method of analysing and simulating the desirable characteristics of existing analogue gain compression equipment has been presented together with a summary of further work in hand at the time of writing. A system for performing these operations in a multi-channel environment has also been described.

REFERENCES

(1) Michael J Kemp: Analysis and Simulation of Non-Linear Audio Processes using Finite Impulse Responses Derived at Multiple Impulse Amplitudes, presented at the 106th AES convention in Munich, Germany, 8 – 11th May 1999. preprint number 4919(J5).