Sampling equalisers

Despite their increased adoption and availability, digital equalisers continue to suffer from a lack of acceptance and endure unfavourable comparisons to analogue EQ. MICHAEL KEMP from Sintefex offers an explanation and proposes a solution.

AUDIO SAMPLING AS a concept in music has been around for some time. From classic analogue instruments like the Mellotron to digital samplers pioneered in the Fairlight CMI and early Apple II-based samplers, the ability to use real sounds transformed the concept of the synthesiser.

At the same time a parallel development in the visual arts saw computer graphic artists adding analogue samples of real objects and surfaces (so called 'texture mapping') to computer generated images to transform the artificial images of the 1970s into the photo-realistic images of today.

In both cases the addition of analogue samples transformed the results obtained. This is mainly because real sounds and images have a complexity that is not easily achieved by digital simulation.

So why not apply the sampling procedure to another area of digital processing that suffers from a similar problem of lack of complexity? This was the question that we posed when considering why typical digital audio processing lacks the sonic complexity of analogue processing.

The background to this question lies in the fact that while digital processing permits us to do many things that were previously not possible, the EQ and dynamics control that is often essential on each track of a recording is not well addressed by its digital equivalent. This, of course, explains the current practice of using a well equipped analogue front-end to digital audio recording: you have to do a lot of the processing in the analogue domain before going to digital if you want acceptable sounds on your hard disk.

However, if the characteristics of analogue equipment can be sampled and applied digitally, maybe there is a way to make digital processing and mixing sound good at last.

So why does an analogue recording chain attract the ear so much? One reason is that each channel of an analogue desk contains EQ and often compressors that have been carefully designed to sound good. Recording engineers also augment these with outboard processors that have unique characteristics that the engineer can choose from to suit exactly the musical component on a recording channel.

With the advent of digital recording the essential elements of EQ and compression have been implemented in the digital processor, often a DAW so as to give as many channels as possible on a single DSP chip. This is done to squeeze as many tracks of audio onto a hardware platform as possible. A typical DSP implementation of an EQ section is shown in Fig. 1 and with minor modifications provides the usual band or shelf boost or cut. The algorithm shown is simple, quick and efficient, and a cheap DSP could probably do several hundred of these operations in a sample period. This might present a user with a multichannel audio mixer at low cost.

This is generally described as the digital equivalent of the analogue circuit shown in Fig. 2 but this is wrong for two substantial reasons: one relates to complexity and the other to the sampling nature of the digital version. Both problems lead to lack of musicality and both can be solved by techniques that I will discuss. First let's look at these two problems.

The analogue representation can never be built with 'perfect' components. Additional amplification elements and interfacing circuits, such as transformers, will be present to complete a real processing device. The inductor will suffer from non-linearities from a variety of sources, for example hysteresis, heating and mechanical effects. In fact, the inductor will often be replaced with active RC networks adding more complexity.

The net effect of the analogue realisation is a complex circuit with a complex effect. The ear will usually detect the insertion of such an equaliser even if it is set to a flat response. A small amount of EQ will result in complex signal processing and in a well-designed circuit the ear will respond favourably to the overall effect.

Contrast this with the digital implementation, which is more or less perfect (more or less because it is possible to program in data errors, for example rounding errors, and these almost always lead to unpleasant effects). To the extent that perfection is achieved the effect on the ear is to hear no effect when the device is programmed to a flat response. When some EQ is applied the ear discerns that the desired effect on the frequency response may have been achieved but does not perceive any of the richness associated with the analogue complexity. The simplicity of the calculations creates no desirable side effect to the EQ. A typical reaction of a user is to continue to increase the amount of EQ applied in an attempt to get a positive effect. The frequent result of a major musical project attempted using such wholly digital EQ is to be over-EQed resulting in a harsh, unpleasant and unsatisfying sound.

Analysis of the frequency response of the digital implementation shows an effect familiar to the DSP designer, that is, the distortion in frequency response as the frequencies processed tend towards the Nyquist limit. This effectively maps an infinite frequency in the analogue system to the Nyquist frequency of the digital system (22.05kHz for a CD system).

In other words, whereas the analogue system response gently tends to its limiting value at infinity (at least as far as the supporting circuitry can manage, which is often well above the audio band), the digital version reaches its limit at 22kHz in an increasing 'crush' – like the bonnet of a car hitting a brick wall (Fig. 3).

Figure 3 illustrates the typical responses of a band lift at 10kHz and 20kHz in an analogue and the 'equivalent' digital EQ circuit. The analogue roll off towards infinity is compressed to the Nyquist frequency in these familiar curves.

This has an important effect on the sound of a simple digital EQ. The ear easily hears the peculiar effects at high frequencies; the rapidly changing gain is perceived as a high Q effect and the effect is often described as a harsh or gritty sound.

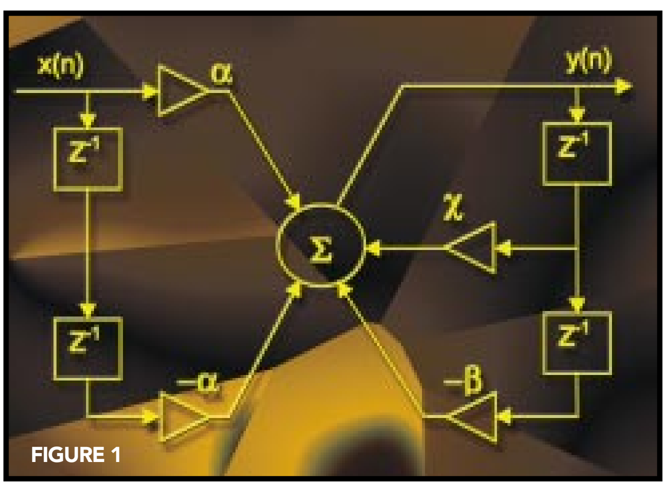

FIGURE 1

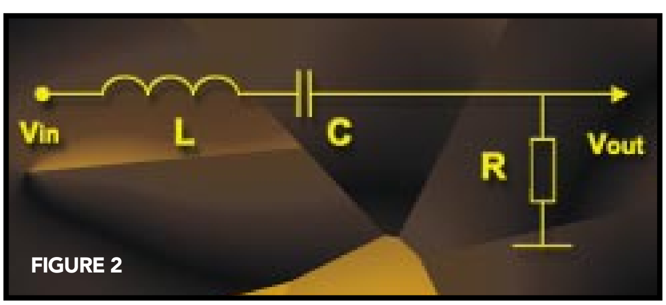

FIGURE 2

The fact that these effects are appalling in the top octave, and continue to be unpleasant an octave or two below, is in my opinion one of the main reasons why there has been a move to higher sample rates. Not because we have suddenly developed the ability to hear into the bat range but because the simple digital EQ is (at least in its higher octaves) basically unmusical and unpleasant, and unfortunately, almost universally used.

To keep this summary of the problems with digital processing short, I will not go into any detail about the limitations of digital dynamics control. The signal path is passed simply through a gain control element that has no significant side effects apart from variable gain reduction. Gain reduction data is determined usually from the incoming audio data stream and variable gain is applied according to simplistic formulae. The resulting gain modulation is unexciting to the ear, and there is no colouration of the signal path. Again, this seeming benefit is in fact perceived by the ear as uninteresting and flat and again the operator is tempted to apply excessive compression further adding to an unexciting final result.

In addition, it is easy, unless careful design is applied, to generate side-bands in the audio due to the modulations that pass beyond the Nyquist limit, generating aliases back into the audio band that sound extremely unpleasant.

None of these problems exist in a typical analogue circuit. Or to be precise, where problems do exist, the analogue processor is rapidly determined to be an unpleasant unit and consigned to the scrap heap. Thus in an almost Darwinian way the older devices that have not been discarded represent a pinnacle of desirable responses. It is not surprising that using them (or emulating them) can be so effective in relieving the digital problems.

A quick note must be added here. There are many fine digital algorithms that add exciting and desirable effects to music, and the power of digital audio editing is a joy to those of us who have had to stick together tiny pieces of magnetic tape to achieve a perfect edit.

Also it is possible to design better EQs and compressors that address some of the problems outlined, but at an expense in processing power that makes such units expensive and rare. However, the best of these units do not in my view address the prime limitation of digital audio, and that is that the effects remain simple. It is from the complexity of analogue processing that much of its musicality derives. However, analogue complexity can be incorporated.

The digital EQ example of Figure 1 (or similar) is an infinite impulse response (IIR) solution that is most often used as it is computationally cheap. A more powerful approach is to change to a finite impulse response (FIR) method. While this is much more computationally intensive, it does offer a path to a genuinely improved result.

The standard FIR and convolution process need not be described here as there are many textbooks on the subject. Instead I'll focus on deriving the values for the FIR table and in extending the method in a novel way.

Many methods exist for deriving terms for a FIR filter to meet arbitrary design rules of filter performance. However, most of these methods have the failing that, because the values are algorithmically generated, the focus is on the production of a given frequency and phase response without reference to acoustically satisfying criteria. As we have discussed above, the ear craves for complexity in a sonic response so the results of such an algorithmic solution is not in itself any better than the results of the IIR method.

Instead, it is a well-known fact that a linear time-invariant (LTI) system can be characterised in frequency and phase response by its impulse response. This suggests that if we have an analogue equaliser that is known to have a good sound it is only necessary to derive its impulse response, then plug these values into the FIR simulator. Once this is done, the digital simulation immediately realises the same frequency and phase response as the original analogue equipment. Listening comparisons show that this immediately sounds convincing to the ear in a way that the algorithmic methods don't.

It is also interesting to note that the processing delay of such a system can approach zero regardless of the length of the FIR used. Further delay is only required when the analogue system also exhibits it as embedded in its impulse response. The delay of the simulation processor is typically less than converting a digital signal to analogue and back to pass through a real analogue unit.

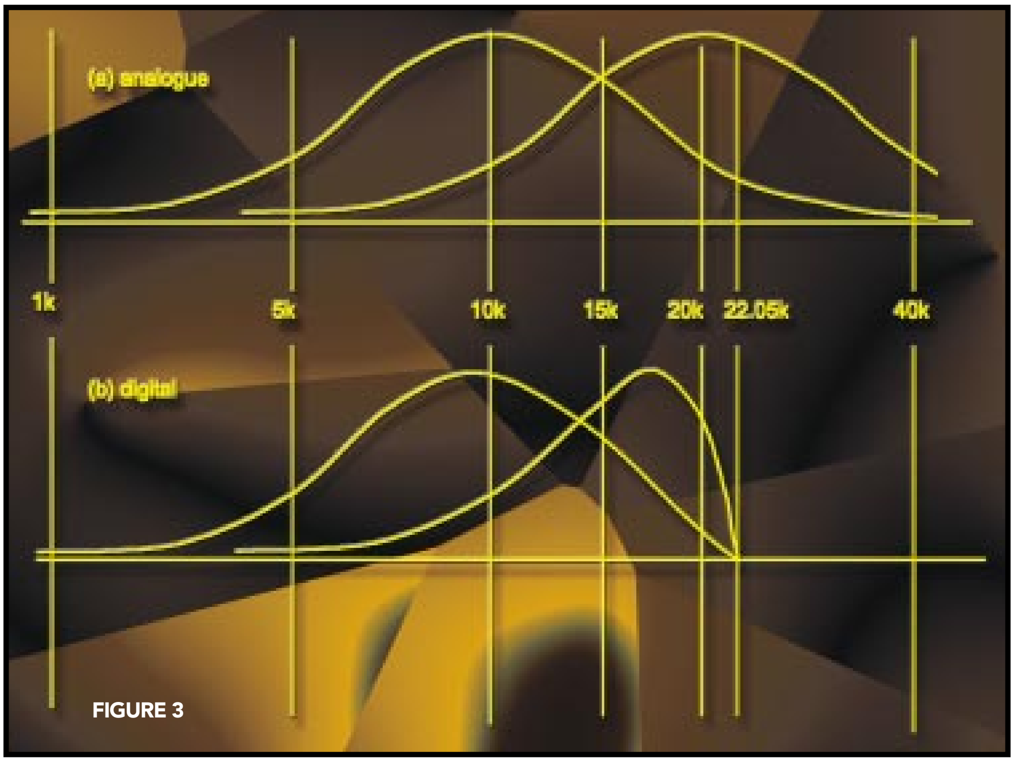

FIGURE 3

In addition, the sample rate crushing effect of the IIR equaliser ceases to be a problem. The original analogue equaliser frequency response is preserved undistorted right up to the Nyquist frequency. No Q distortion occurs and the clarity of HF equalisation is immediately evident. Figure 4 shows the analogue and corresponding digital responses.

However, the simulation still falls short of the original analogue device. This is because the method described can only simulate an LTI system, and unfortunately an analogue equaliser has, as part of its complexity, non-linearities that are habitually exploited by recording engineers. It is a well-known practice to drive an analogue processor harder to change its sound. This applies to virtually all analogue processors, and especially to older devices with vacuum tube and inductive (or magnetic) components, and to any system with an analogue gain control element. Capturing this variation is the goal of the improved convolution method now described, which adds the additional complexity of level sensitivity to complete the analogue simulations and give the ear the processing effects it craves.

In this modification of the convolution process the original equipment now has its impulse response determined at a whole set of levels, typically from the peak excitation possible in the analogue device down to about 40dB below this value. Beyond this it is difficult to measure the response due to noise in the system and the reasonable assumption is made that such low-level responses are now truly linear.

The convolution algorithm is now adapted so that the incoming signal level is continuously assessed and an appropriate response is selected from the set obtained in the original sampling process. This response is applied and results in the originally sampled impulse response for that level being applied proportionally to that sample, which is then linearly superimposed on each subsequent sample.

In practice it is necessary to linearly interpolate

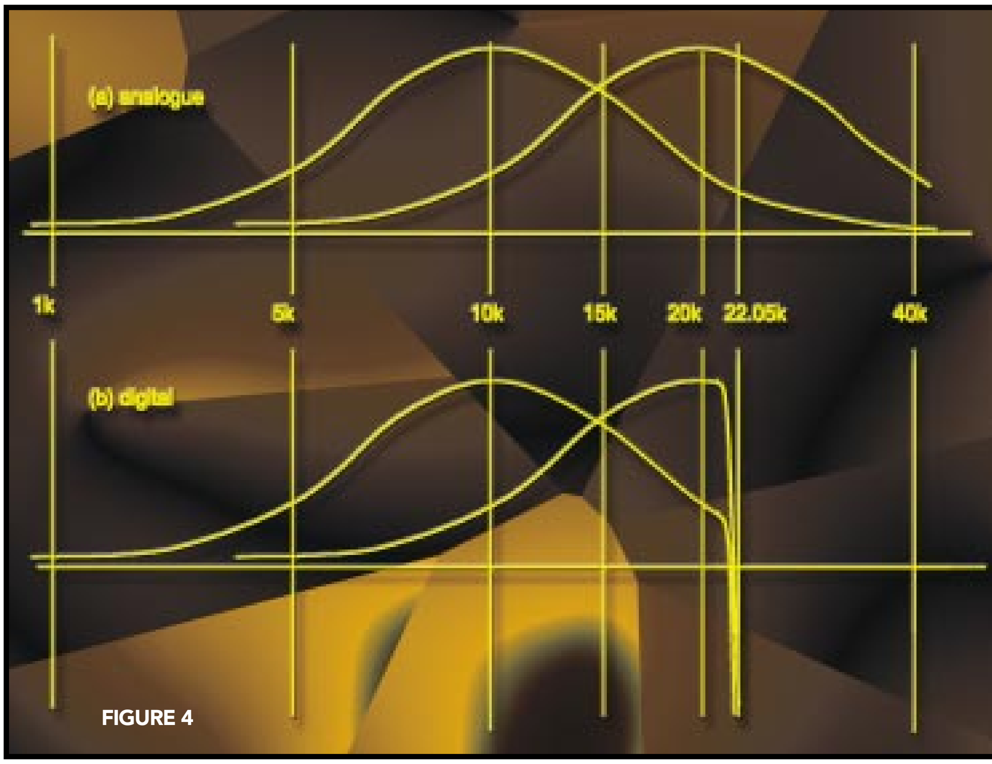

FIGURE 4

responses from either side of any given input level signal as, for example, in a 24-bit system there are 223 possible input signal levels (more in cases where the level determination requires more than one sample to be assessed) but, of course, fewer samples are available from the original equipment impulse response set.

This simulation technique now performs remarkably similarly to the original analogue equipment. Frequency response, phase response and change in response with signal level are replicated. Signal paths of arbitrary frequency response may be sampled in the few minutes it takes to pass the impulse test signal through the device.

The added complexity and memory requirements of dynamic convolution require substantially more DSP power than the simple IIR system of Figure 1. However, it now appears that a fact well known in analogue design also applies to digital design. A cheap and nasty algorithm does not sound good any more than a cheap and nasty analogue implementation. It is time to pay for better sound. Digital audio will only start sounding good when the digital processing starts to get as complex as analogue processing has always been.

Luckily with the continually increasing performance to price ratio of DSP devices the actual cost is not prohibitive. We have found that a single channel implementation of a 2048 step dynamic convolution algorithm (where the linear interpolation implies two parallel 2048 convolutions are performed) can be achieved in a pipeline of four SHARC DSPs each fitted with 1Mword of 32-bit impulse memory.

Dynamic convolution models a static channel of analogue signal processing. However, it is also possible to make multiple samples of a variable signal path and perform an interpolated real-time change in the simulation. In this way it is possible to store many samples from a device, such as a variable EQ unit, and to provide user controls to match variations of a virtual control panel to the original control settings on a sampled device.

As an example from our Sintefex Audio library of sampled devices, each sampled EQ embodies many hundreds of sample sets containing many tens of thousands of numbers. (In fact some EQ samples contain over a million 32-bit floating point values to define their sound.) This is well contrasted with the three numbers required to define the digital EQ of Figure 1. It is hardly surprising that such simplicity cannot compete in sonic quality with the sampled simulation.

It is also possible to apply the sample technique to dynamics processors with equally successful results, but this is perhaps the subject for another article.

For equalisation and dynamics control, once sampled, all the benefits of digital processing follow. As an example, we have built a library of favourite processes, so a rack full of analogue outboard can reside in the memory of the system ready for recall at the touch of a button. It is possible to choose just the right effect by trying each of the appropriate processors quickly, and, once the correct process has been chosen the details can be stored for exact recall again and again, as often as necessary. It is now possible to apply the same processing identically to two channels for perfectly matched stereo, or across all the channels of a 5.1, 6.1 or 7.1 mix – impossible with true analogue gear. For dynamics, gain matching is perfect, side chain EQ is always available, and the limitations of attack and release times of many original models can be relaxed.

Finally, the hum and noise of analogue equipment is a thing of the past. Analogue sounding effects can be applied with the full digital dynamic range, and silence follows as the signal fades. There is now no fear that, at the vital moment, a vacuum tube will fail and take with it its unique sound.

Contact

SINTEFEX AUDIO, PORTUGAL:
Tel: +351 282 361748
Website: www.sintefex.com

More reading

- **Implementing IIR/FIR Filters with Motorola's DSP56000/DSP56001 by John Lane and Garth Hillman, Digital Signal Processing Division, published by Motorola Inc. 1993.**
- **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, May 1999. Preprint number 4919(J5).**
- **Michael J Kemp: Analysis And Simulation Of Analogue Dynamic Compressors And Limiters In The Digital Domain, presented at the 109th AES convention in Los Angeles, CA, September 2000. Preprint number 5185.**