

The Lost Art Of Sampling: Part 1

By Steve Howell



Most modern musicians use samples, even if only in S&S keyboards or virtual instruments. But sampling itself has become something of a lost art. In the first of a short series on rediscovering this skill, we look back at how the technique and the technology developed.

Sampling technology has become so widespread that it is no longer considered remarkable. Indeed, in many ways, it has become 'invisible' — so widely accepted and taken for granted that no-one notices it any more. Everywhere you look, you'll see sampling in action: in the stored messages on modern ansaphones, in those muffled train and airport announcements you can never quite catch the gist of when you're in a hurry, and in those dreadful menu-driven customer service lines so beloved of companies claiming to offer "a better, more focused customer service experience". Wherever we turn, we have first-hand experience of 'sampling' in one form or another.

It's the same in the modern music-making industry — sampling technology is at work in the vast majority of synth, keyboard and virtual-instrument products on the market. Whereas synths were once all powered by analogue oscillators which generated a limited range of electronic waveforms, these circuits have largely been replaced in modern hardware synths by chips containing samples of a colossal range of instruments, including analogue-synth-style sawtooth, square and triangle waveforms to complete the illusion. Even some of the 'modelled' analogue synths use carefully engineered multisamples of analogue (and other) waveforms as the basis of their synthesis methods.

But even though samples lie at the heart of much of what we do, few of us actually sample audio ourselves any more. So-called software samplers also rarely sample, although they're happy to sort raw audio into multisamples and keymaps if you can find some other way of getting the audio into your computer in the first place. And nearly all sample libraries are supplied on CD-ROM these days, or as sample libraries with a virtual-instrument 'wrapper', which completely side-steps the former need to sample, trim and if necessary loop audio before we can use it. Not surprisingly, outside the select world of the sound designer, there are few people who can be bothered these days.

This is all very well if you remember the days of searching for seamless sample looping points with horror, and can't get enough of purchasing your samples in a ready-to-go, easy-to-use format, but what of all those people who've come to the world of music technology in the last five years, and never learnt any sampling skills in the first place? If this applies to you, or if you're sick of spending soul-destroying afternoons trawling through your vast sample libraries looking for 'the right sound' and would prefer to be able to create your own from scratch, or if you'd just like a refresher course in what sampling is, and what you need to know about it in order to handle it yourself, then this series is for you.

In The Beginning...

Let's start at the very beginning – what *is* sampling? Strictly speaking, it means 'digital recording and



The groundbreaking Fairlight CMI

playback' – the act of turning a sound source into a stream of 0s and 1s that a digital processor can deal with.

The term was first coined in the late '70s by Kim Ryrie and Peter Vogel, inventors of the Fairlight CMI, to succinctly describe one of the features of their then-new digital synthesizer (for more on the Fairlight, check out our Retrozone feature in *SOS April 1999*).

As with some of mankind's greatest inventions, sampling was originally a kind of 'bonus feature' that was included almost as an afterthought. The Fairlight was designed as a digital synthesizer on which waveforms could be constructed using additive synthesis, or 'drawn' using the innovative light pen and touchscreen. It was also possible to define start and end waveforms and 'morph' between the two to create dynamic tonal movement throughout the course of a note. Of course, the Fairlight also had a sequencer, and as such, was also the world's first 'workstation'. Yet it proved to be the sampling feature that captured the imagination of the musicians and producers who were wealthy enough to slap down the £20,000 required to buy a Fairlight CMI in 1980. In retrospect, perhaps this is not so surprising – it was already possible to sequence electronic synths, and the idea of being able to do the same with the sound of *acoustic* instruments was a powerfully attractive concept.

However, not everyone understood the great leap forward at first – another recurring theme in the story of many of mankind's greatest discoveries! One long-defunct music technology magazine of the day sneered at the new 'sampling' feature, wondering why anyone would want to use such a powerful digital synth in this way, and concluding that it was a gimmick. In fairness, even Ryrie and Vogel regarded their invention as a digital synth first and foremost, and whilst the sampling feature allowed complex sounds to be recorded and played back, it gave you very little control over the sound other than the pitch and basic envelope, whereas the digital synthesis components gave precise control over almost every aspect of a sound. But as time went on, it became clear that people weren't buying the Fairlight for these innovations – it was the sampling feature they wanted, and the technique was to change the world and the music it made... forever!

Of course, even with its headache-inducing price tag, the sampling facilities on the Fairlight were very limited, offering recording times of a few seconds at best (and usually less, as sample RAM was then unbelievably expensive), and all at a very poor quality compared with what we're accustomed to today. A 'looping' facility was added to allow a portion of the recording to play back over and over again for as long as a note was held, but that was all. At the time, however, it was a revolution. The general idea of using audio recordings within larger recordings was not new, and nor was the more specific idea of using short

recordings of real instruments playing individual notes to create melodies — both of these ideas had been explored using analogue technology from the 1940s onwards (see the 'From *Musique Concrète* To The Fairlight' box), but the digital nature of the Fairlight CMI, and its built-in sequencer, made the concept much more workable and immediate.

Inevitably, the success of the Fairlight meant that other companies began making their own samplers (see the 'From Fairlight To Expansion' box), and over the years, stiff competition in the field ensured that the specifications of samplers improved out of all recognition, prices dropped drastically, and eventually, we arrived at the situation we have today, where broadly speaking, what comes out of a modern sampler is pretty much what went in. It wasn't always that way, but today's technology means that you don't have to think too much (if at all) about the mechanics of the process.

Sampling Dissected

Those of you unfamiliar with what goes on beneath the user interfaces of modern samplers are no doubt wondering how exactly the process *does* work. Well, if you remember that sampling is basically digital recording, then in order to begin to understand sampling, you need to understand the principles of digital recording.

Analogue signals, like the sounds our ears pick up, are continuously variable, whereas digital signals are made up of strings of zeros and ones so that a computer can process them. To convert analogue audio into digital data, we need to feed it through an analogue-to-digital converter (or A-D converter, also known as an ADC). This slices the audio waveform up into little sections, each of which is assigned an amplitude (loudness) value which can in turn be expressed as a binary number, so that a digital audio processor can understand it, or a computer can store it. Each of the loudness values taken, properly speaking, is called a sample, so this process is known as sampling the waveform. To replay the sound, the binary amplitude values are fed through a digital-to-analogue converter (or D-A converter, or DAC) which reconstitutes the slices at different loudnesses, and puts them back together in the right order as a continuously variable signal that can feed your speakers, and which you can then hear as a sound.

If this explanation has left you totally lost, try thinking of cinema film — I often find that making this analogy helps people to grasp the concept. As you probably know, cinema film is a strip of still photographs which, when replayed fast enough, gives the illusion of continuous movement. Typically, when making a film, 24 frames (or still photos) are taken every second. The more frames there are in any given second, the smoother the movement on screen will be. 24 frames per second (or fps) is largely accepted as the optimum rate for commercial cinema film.

The same is pretty much true of digital audio. When you record a sound digitally, you are taking a series of 'stills' of the loudness of the continuously variable input signal, each of which is then stored as a binary value in Random Access memory (RAM) or, these days, on hard disk. Then, when you play back the sound, the stored loudness values are played back in the correct order, reconstituted by the DAC as a variable waveform, and what we hear is an accurate replica of the sound we recorded. That's the idea, anyway!

From Fairlight To FXpansion

The Fairlight was the hippest piece of music technology for years, but its high price tag meant that it was also very exclusive. Of course, this may have only added to the buzz it generated! For the price of a Fairlight CMI on its release, you could buy a country house in the UK, so only the wealthiest of musicians and producers could buy one. Inevitably, a race began to get the price of the technology down to a more affordable level. New England Digital expanded their Synclavier digital synth to include sampling following the success of the Fairlight, but this was equally expensive.

Emu Systems (who, until this time, were primarily manufacturers of modular analogue synths) were the first to truly take things further with their imaginatively named Emulator. This offered basic 'sampling' features similar to those of the Fairlight in a £5000 wedge-shaped keyboard. They went more upmarket with their £9000 Emulator II (below), which added filters, envelopes and other synth functions, as well as individual outputs, timecode synchronisation and other features into the equation. But both of these were still a bit pricey for most people.



E-mu's Emulator II

In the meantime, while mere mortals waited for the price of sampling technology to drop, some were exploiting the new breed of digital delay lines (DDLs) to 'freeze' (ie. sample) audio and then trigger it. The downside of this technique was that sampled sounds could not be 'played' (although it was possible to tune the sample up and down). The DDLs didn't have any form of storage, either, and so the sound was lost as soon as the unit was switched off. Despite these limitations, and in the absence of anything else, many producers and engineers used DDLs to trigger drum sounds. They also used them to sample, say, backing vocals – they would mix down the BVs into the DDL and then 'spin them in' wherever they were required in the song.

As with anything new, the first DDLs were expensive... very expensive! Probably the first manufacturer to offer this facility was UK company AMS with their DMX DDLs. They were initially mono, although the company later made a stereo version, as well as one that could be played from a keyboard that generated control voltages, although it was several thousand pounds in the UK. But the DMXs had a unique feature for the time – expandable memory – and many bought them so that they could 'spin in' even longer sections of audio. Producer Steve Levine had his AMS fully expanded (at a cost of several thousand UK pounds) when he was producing the band Culture Club in the mid-'80s. Another UK company, BEL, also produced a similar DDL, but this was considerably less expensive, and consequently, the distinctive blue BD80 DDL graced the racks of many studios at the time.

Of course, it was only matter of time before the cost of the technology fell, and in 1983, aided by mass-production techniques, Boss brought out their DE200. This was a flexible DDL in its own right, with modulation for chorus and flanging effects, but it also had (low-quality) sampling, and all for a few hundred quid. This was followed in the mid-'80s by Ensoniq's Mirage, which, at £1200 in the UK, went head-to-head with the S612, the first sampling product from another then-unknown company, Akai. Sequential Circuits joined the fray shortly afterwards with their Prophet 2000 and 3000, but sadly went bust before their sampling products could reach maturity.



Ensoniq's £1200 Mirage sampler set an all-time low price for sampling keyboards in 1985

The Mirage was a landmark product at the price, but it sounded pretty rough and was a pig to use, with just a two-digit display. It also had a limited sampling memory of just 144K! Akai's S612, on the other hand, offered much higher-quality sampling and longer sampling times.

The mass market really woke up to the potential of sampling following the release of Akai's S900 in 1986. It offered good audio fidelity, up to a minute of recording (albeit at compromised quality), individual outs and advanced sample- and program-editing facilities for the time. It was also rackmount and fitted in perfectly with the emerging world of MIDI. The S900 faced stiff competition on many fronts, as everybody was beginning to make their own samplers, from Roland with their S-series (S50 keyboard and S550 rack) to Yamaha with their TX16W and even Casio, who offered the FZ1. Ensoniq also followed the Mirage with their popular EPS series, which incorporated workstation-like sequencing features. But none of these could really compete with the market supremacy of the S900. In fact, the only real competition came from Emu's Emax, which offered most of the features of the big Emulator II, but at a fraction of the cost..



The former 'gold standard' of hardware samplers, the Akai S1000

Of course, the S900 paved the way for the CD-quality S1000 and coincidentally, Emu released their similarly endowed EIII at around the same time. From that moment on, it was pretty much a battle between the two giants of sampling, Akai and Emu. Roland hung on in there with the S330 and their new S700 series, but by the early '90s, everything had come down to Akai and Emu, who seemed to operate almost in parallel, each one offering comparable products but with slightly different features and functions. And so allegiances were born, and users of each brand swore an evangelistic commitment to their chosen sampler that rivalled the passions in today's Mac vs PC debates. The final result was probably a draw – the US just seemed to prefer Emu samplers, whilst Europe favoured Akais.

And then, at the turn of the Millennium, software samplers arrived, and the whole Akai vs Emu debate rapidly became academic. As computing power had increased on both the Mac and PC, it had become increasingly viable to make a software sampler that could give the hardware equivalents a good run for their money. Akai and Emu continued to bring hardware samplers to the market, but unfortunately, the writing was on the wall. The software versions were not necessarily better, but they were cheaper and integrated more easily with users' preferred DAW and/or sequencer. Piracy also played an enormous part in the downfall of the hardware sampler – whereas before, aspiring musos would invest in a budget Akai S2000 or Emu ESI32 with every intention of upgrading to an S5/6000 or E6400 Turbo when budget allowed, those same people could now get a knocked-off copy of a software sampler from a friend or an Internet-based peer-to-peer network for nothing.

Many people (myself included) still prefer hardware samplers for a variety of reasons, including sound quality, reliability, portability, and the lack of inexplicable software conflicts, latency, and CPU strain, but from a standpoint halfway through the first decade of the 21st century, it is clear that hardware samplers have had their commercial day, and the future of sampling – for the moment, at least – lies with software.

Sample Rates

Sticking with the film analogy for the moment, let's move on to consider the idea of sampling rate. As I've mentioned, it is accepted that 24 frames per second is the optimum number to give an accurate perception of smooth movement when played back through a film projector. With CD-quality audio, it is largely accepted that 44,100 audio slices (or samples) per second is the optimum number to give an accurate perception of the original sound. This figure, expressed as the number of samples taken per second, is known as the sampling rate. 44,100 samples per second can be written as 44,100Hz (Hertz being the standard unit of frequency, expressed per second), or, as it is more commonly written, 44.1kHz. But why such an odd figure?

Well... the highest frequency a human being can hear is agreed to be around 20kHz (20,000 cycles per second). As it happens, this is usually untrue – at best, it's more likely to be 17 or 18kHz, if that! Newborn babies might be able to hear frequencies as high as 20kHz, but as we get older – or attend too many deafening concerts – the highest frequency we can hear tends to drop. To accurately capture any frequency in a digital audio system, you have to sample it at twice that frequency. Therefore, to capture potentially audible frequencies of up to 20kHz, you need a sampling rate of at least 40kHz. This principle was discovered by one Harold Nyquist back in the 1950s when he was working at Bell Laboratories and researching the possibility of digital audio transmission for telephone systems. It is known as the Nyquist Theory. To explain why the sample rate must be twice that of the upper frequency limit, I'm afraid we'll have to return to the film analogy!

No doubt you've seen old, silent movies – Charlie Chaplin, The Keystone Cops, that kind of thing – and the movement seems jerky. This is because in the early days of cinema, they used slower frame rates – typically 16 frames per second, sometimes less. It was thought that this would be sufficient to create the illusion of smooth movement on screen, but as it turned out, it wasn't. Increasing the frame rate made the illusion believable, which is how the film industry arrived at its rate of 24fps.

It's similar with audio. At lower sample rates, the sampled audio is not a true representation of the original analogue input signal, and is 'jerky' – which in audio terms, translates into 'fuzzy' or 'murky-sounding' audio. But there's another reason why high sample rates are required to play back audio accurately. Back to cinema film again!

Aliasing

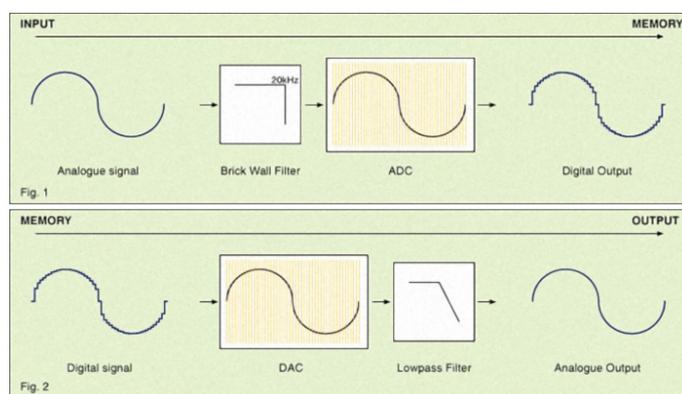
You must also have seen films where a moving vehicle's wheels appear to be going in the opposite direction to the vehicle itself. If the wheel is spinning faster than 12 revolutions per second (ie. faster than half the film's 'sampling rate' of 24fps), the wheel completes more than half a revolution per frame, and our brains are then unable to tell whether it got like that by turning more than half a revolution forwards, or by turning less than half a revolution *backwards*. So, for example, if the wheel is spinning at 18 revolutions per second, it will *appear* to be going backwards by six revolutions a second.

The reason for this is that the camera (and our eyes) cannot distinguish between positive and negative frequencies – to the camera, a revolution of 18 24ths is exactly the same as a negative revolution of six 24ths. This phenomenon is known as 'aliasing'.

The same is true of sampled audio. Any frequency that exceeds the Nyquist frequency (ie. half the sampling rate) will be 'reflected back' into the audio spectrum by the amount that exceeds the Nyquist frequency. For example, if the sampling rate is 20kHz (giving a maximum upper frequency response of 10kHz according to the Nyquist

theory) and the sound being sampled (for example, a cymbal) has a high-frequency overtone at 12kHz, that overtone will be 'wrapped around' at 8kHz. 8kHz is the Nyquist frequency of 10kHz minus 2kHz, the amount by which the cymbal overtone of 12kHz exceeds the Nyquist frequency. This undesired 8kHz overtone (also known as an alias) is unlikely to be mathematically related to the frequencies in the original sound, which is a prerequisite for sounding pleasingly harmonious to our ears. So the chances are that it will sound highly enharmonic – in a word, nasty! Cymbal sounds in particular consist of many complex frequencies, and are likely to have other overtones at high frequencies which will also be reflected back into the audible spectrum in mathematically unrelated ways, creating further sonic mayhem. You can now begin to understand why early and/or cheap samplers, with their low sample rates, sounded pretty horrible. There was a whole lot of aliasing going on!

So surely a sample rate of 40kHz (and a consequent theoretical upper frequency response of 20kHz, the accepted practical upper limit of human hearing) will guarantee that this will not happen. Or will it? The fact is that many natural sounds *do* contain frequencies in excess of 20kHz. We might not be able to hear them, but that doesn't mean that they're not there. And what we would certainly hear are the aliases from audio frequencies beyond our hearing, which would be reflected back into our audible range around the Nyquist frequency in the manner I've described, even when the sampling rate is 44.1kHz.



A simplified diagram of the sampling/digitisation process.

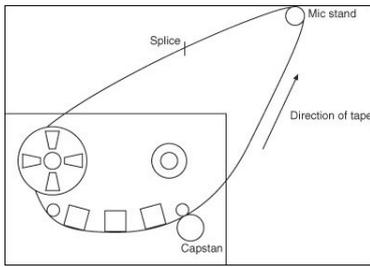
The reason we don't hear these aliases in CD-quality digital audio is that modern ADCs filter the audio input to prevent any frequencies above 20kHz entering the system. These are known as 'brick-wall' filters because of the steep shape of their frequency responses, and they rather brutally remove anything above 20kHz to exclude the possibility of audible aliasing making it into digital recordings. You can see the complete process in the diagram on the left. The analogue signal passes through a brick-wall filter, where any frequency above 20kHz is removed. That signal is then 'scanned' and sliced at the sampling rate in the ADC and the resulting digital signal is stored in memory (in RAM or on disk). To play the sound back, the digital signal passes through a DAC where the slices are reconstituted but, because there might be some 'rough edges' in this process (which would be heard as distortion), the signal passes through a low-pass filter to smooth out the signal, and the reconstituted analogue signal is passed to the outputs.

Armed with this knowledge, why is a sample rate of 44.1kHz required, given that to reproduce upper frequencies of 20kHz, you only really need a sample rate of 40kHz? Well, some overhead was added to allow for the various filtering processes involved, and the Audio Engineering Society settled on 44.1kHz as the *de facto* standard for CD-quality recordings in 1985.

Of course, this description is highly simplified, but is pretty much all (if not more than) you really need to know without resorting to mathematics. If you do want to find out more on the intricacies of digital recording, refer to Hugh Robjohns' excellent series on the subject, which appeared in *SOS May to October 1998*.

From *Musique Concrète* To The Fairlight

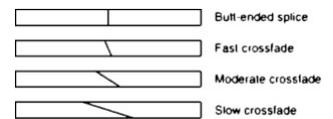
Although the digital approach to sampling introduced by the Fairlight made the process of using recorded sound in larger compositions much easier, the technique itself considerably pre-dated the late 1970s. As far



A diagrammatic representation of how lengthy tape loops used to be set up 'in the old days'.

back as the mid-1940s, Frenchman Pierre Schaeffer was pioneering a new musical form called *musique concrète* which involved recording acoustic sounds to short lengths of recording tape, splicing them together, speeding them up and slowing them down, reversing them and otherwise manipulating them to create collages of sounds. His first piece (in collaboration with fellow French composer, Pierre Henry) was *Symphonie Pour Un Homme Seul* and was made up entirely of sounds from the human body. Later works included sounds from locomotives to kitchen utensils.

The technique also used tape loops. Before the age of digitised audio, the only way to have a sound repeat endlessly was to splice a tape recording end-to-end. In *musique concrète* it was common practice to make the tape splices at various angles to create smoother edits, a technique which later made its way into mainstream tape-based recording techniques. With a right-angled splice, the transition from one tape section to another is of course instant, but by using varying splice angles, the recordist could smooth out the transition. Of course, nowadays, such things are handled effortlessly in any decent audio-editing software with a single mouse drag, but before the age of digitised audio, tape splicing was an art form that required considerable dexterity and skill that could only be acquired through experience.



'Old-school' crossfading: using different angles of tape splice

Once the section to be looped was spliced together, it needed to be run through a tape machine, and if it was a long section which wouldn't fit inside a reel-to-reel recorder, you had to resort to all sorts of Heath-Robinson solutions, including the famous one of wrapping the loop around a spare microphone stand (see the diagram above). This was a hit-and-miss affair, as the tension of the tape had to be carefully adjusted – if it was too tight, you ran the risk of scraping the oxide off the tape (as well as potentially introducing wow and flutter), and if it was too loose, the tape could slip and be chewed by the transport mechanism.

Assuming all this could be set up correctly, a tape section could be made to repeat endlessly. Schaeffer even invented a unique tape recorder called the Phonogene that allowed him to play loops at 12 different pitches from a keyboard, with 12 carefully constructed capstan rollers which (much like gears on a bicycle) could be selected from the keyboard. The Phonogene also had a two-speed motor for octave transposition.

Musique concrète wasn't very popular with the traditional classical music fraternity; the notion of using acoustic sounds in a non-real-time performance was not classed as 'music'. However, that didn't stop composers such as John Cage, Edgar Varèse, Karl-Heinz Stockhausen and Iannis Xenakis experimenting with the form. By the 1950s, Louis and Bebe Barron were forging ahead, using new techniques to create the world's first totally electronic film soundtrack for the classic sci-fi film *Forbidden Planet*. With the aid of unique electronic circuits the couple developed themselves, the soundtrack was recorded employing many of the techniques used in *musique concrète* (splicing, reversing, looping, and so on) to create the stunning sonic backdrop to the film (credited as 'Electronic Tonalities').

Back in the UK, from the mid-'50s onwards, the BBC's Radiophonic Workshop were responsible for bringing these techniques more into the mainstream. The Workshop's purpose was simply to create music and audio accompaniment for BBC programmes. But within that brief, the composers working there exploited the new

techniques, including tape splicing and primitive synthesis. If there was one TV programme that catapulted these innovations into everyone's living rooms, it was *Doctor Who*. Although it sounded like the product of a synthesizer and a multitrack recorder, the original 1963 version of the theme tune used neither – it was instead assembled using a combination of tape-splicing, *musique-concrète*, and sound-on-sound techniques. The source sounds were either 'found' recordings or the output from electronic signal generators, processed via nothing more complex than simple tape delays and then painstakingly layered by repeatedly bouncing from tape recorder to tape recorder, adding layers each time and syncing the various recorders involved purely by hand and ear, as the BBC had no proper multitrack facilities at the time.

Around this time, the BBC took interest in a new product that had just come out – the Mellotron. Often erroneously described as the world's first sampler, this was in fact a playback-only tape-based instrument, and it wasn't even the first of those. The Mellotron in turn derived from another tape-based instrument invented as far back as the late '40s – the Chamberlin. The full story can be found in the review of the new Mellotron in *SOS August 2002*, but here it will suffice to say that its inventor, Harry Chamberlin (a keen organist and home tape recordist) supposedly had the idea when recording himself playing the organ. The story goes that he realised that by making one-note recordings of real instruments and triggering the appropriate ones from a musical keyboard, he could create a keyboard instrument with the potential to play back any sound – an unheard-of idea at the time.



The story of the development of the American Chamberlin into the UK-based Mellotron is confusing, and most versions are hotly disputed, but it is safe to say that the design of the former inspired the more reliable latter! Not surprisingly, the more robust Mellotron proved more popular in the end, and is now revered as a classic instrument, perhaps due to the many influential bands that used it. These included the Beatles, the Kinks, and the Moody Blues in the '60s, and later King Crimson, Yes, Genesis, and Tangerine Dream. More recent users include the eels and Oasis.

However, the Mellotron had potential for use as more than just a musical instrument. In the early '60s, someone at the BBC had the idea of using Mellotrons as sound-effects generators with different effects on each key, and so several were made especially for the Beeb to add sound effects to TV and radio programmes with. From the late '60s onwards, synthesizers, which were at first marketed as being able to create any sound electronically, began to make inroads into the Mellotron's domination. However, the early synths were all monophonic, and so the Mellotron remained popular wherever something capable of playing chords was



The very rare Birotron (seen at the now-defunct UK Museum of Synthesizer Technology in the mid-'90s)

needed. Gradually, though, the frustrations of using them put people off. They were very heavy and could be unreliable. What's more, the tape-based instrument recordings were not looped – so they lasted a finite eight seconds (at best) and then stopped, often very ungracefully! Rick Wakeman (of Yes, and also the session Mellotron player on Bowie's 'Space Oddity') was so frustrated by this limitation that he funded the development of Dave Biro's Birotron. This used continuous tape loops that rotated constantly – when a note was played, the tape was simply pressed against a tape head and playback began at any arbitrary point in the loop. To overcome the wow and flutter and slew that was inevitable when the rotating tape loop made contact with the head, each key had a simple attack parameter to soften the attack and hide the artefacts. But the Birotron wasn't a success, and only a handful were ever made.

However, by this time, playback instruments were being developed in the most unlikely places. Toy manufacturer Mattel – known best for their Barbie doll – were responsible for the Optigan. This used optical



Mattel's Optigan. This one was photographed in Damon Albarn's studio in the late '90s

disks to play back pre-recorded instrumental sounds, drum patterns and musical accompaniments. The discs had waveforms on them which were read by a light-bar reader inside the machine and then turned into sound, much like the old optical soundtracks used on film. Like the Chamberlin and Mellotron before it, the Optigan's principle market was home organists, and, by buying pre-recorded disks, you could 'acquire' accompaniment in different styles to play along with. In many ways, the Optigan pre-empted the various 'arranger' keyboards of today. But like the Chamberlin

and Mellotron before it, the technology used was embryonic and unreliable. The disks were scratchy and the sound quality wasn't good, and as such, the Optigan was not a success.

By the late '70s, the advent of the polyphonic synthesizer pretty much spelled the end of the analogue playback keyboards. The original tape-splicing techniques had also largely been abandoned as being too time-consuming – just a few minutes of *musique concrète* could take weeks of work or more. The last creative and commercial hurrah as far as the old tape-looping techniques were concerned was arguably UK band 10cc's 'I'm Not In Love', in 1975 (the detailed subject of SOS's Classic Tracks June 2005).

What is evident throughout this story is that musicians were obviously keen on the idea of capturing and manipulating acoustic sounds, no matter how awkward or time-consuming the methods required. It was no wonder that when the Fairlight appeared, its sampling function was embraced so warmly.

Higher Rates

Of course, times have moved on since the development of the CD-quality specification in the mid-'80s, and we're now looking at possible rates of 96kHz or even 192kHz for digital recording, which result in ever higher upper frequency responses of over 40kHz and over 90kHz respectively. This fits with some audiophile schools of thought, which maintain that whilst they are not audible in the conventional manner, sounds over 20kHz can nevertheless be 'felt' or perceived in some way. However, we do have to look at this within the context and perspective of current audio recording and reproduction technology.

Most mics can't record frequencies much above 20kHz (maybe 22kHz at best) and most amps and speakers cannot reproduce frequencies much above these figures either. Some pro monitor speakers can extend up to 40 or 50kHz, as can specialised amps, but your average amp/speaker isn't going to get close! Also, there are very few A-D and D-A converters that can accurately handle these rates either unless you start looking at specialised (and expensive) outboard converters. And then there's the fact that many conventional musical sounds don't have frequencies that come anywhere near even the 20kHz upper frequency limit of the CD standard. Knowing this, it starts to seem a bit silly when you see processing power and extra storage space being eaten up to record at these supersonic sample rates – particularly for kick drums and basses. Even sounds such as piano, strings, guitar, and most drums and percussion can benefit little from being recorded

at these rates – especially when you factor in the technical limitations of analogue recording and playback mechanisms, such as those found in even modern mics and monitors.

Nevertheless, there is no doubt that there are many in the modern recording community who swear that they have (say) Fender Precision bass samples recorded at 96kHz that sound superior. For the record, my opinion is that once said Fender bass (or whatever) is buried in a track with other instruments, any hypothetical benefits the ultra-high sampling rates might bring are largely going to be lost – especially when the track is mixed down for use on a 44.1kHz CD and played through the average hi-fi... or worse, made into an MP3 and listened to through iPod earphones on a tube train!

However, I don't wish to start sounding too much like an instalment of Grumpy Old Men, so let's summarise. In practice, 44.1kHz is more than enough to adequately sample most instruments for most musical (and non-musical) application for playback on most systems. If you want to sample at higher frequencies, that is, of course, your decision if you have the equipment to do so, although it will of course make your sampler work twice as hard (or more) to achieve whatever sonic improvements it does. Certainly the polyphony will be restricted at the higher rates in hardware samplers, and in software samplers, the host CPU will have to work harder, which will either result in the same restrictions or possibly in worse, more intrusive problems, such as dropouts, clicks or outright crashes.