Sony DRE-S777 Sampling Reverb

A sampling reverb? It's not what you think. Brad Watts investigates a fascinating new concept in reverberation.

Sony are one of the big players in the audio world so I suspect they can sink obscene amounts of money into developing new angles in professional audio devices. One only has to see a Sony Oxford digital recording console to realise how dedicated Sony are to pro audio. Everything from MiniDisc players to multitrack recorders wander out of Sony's factories, and have done for many years now. (I wonder how much Sony's interest in the artist and repertoire side of the recording industry affects their professional recording products?)

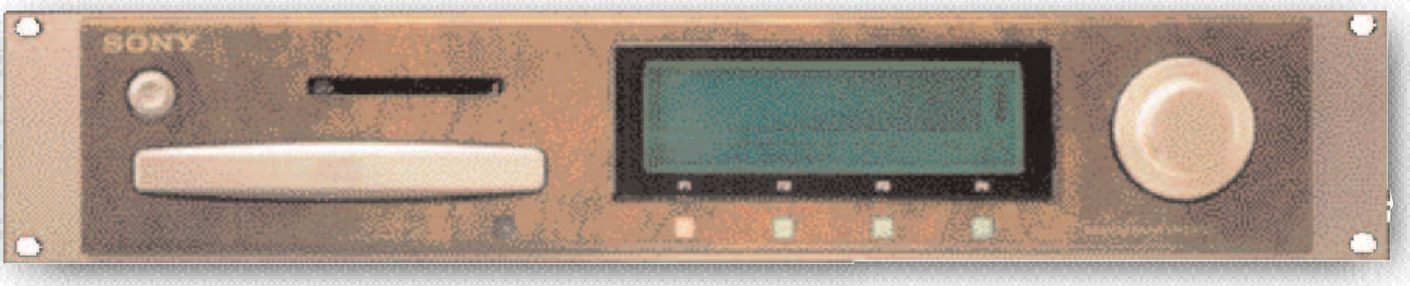

With their newest product, the DRE-S777 Sampling Reverb, Sony have combined a number of technologies to realise a new direction in creating artificial reverb spaces (or should I say recreating real reverb spaces?). Sony generously lent AudioTechnology one of only a handful of these units currently available in the world, so we could give our impressions of this new reverberation method.

When I first heard of this 'sampling reverb', I must admit, I was slightly sceptical. I was thinking "oh yeah right… as if any recording engineer or studio person is going to have the equipment and room to make it worthwhile sampling their own reverbs". And what use is a sample of a reverb, anyway?

My initial visions didn't prepare me for what this 'sampling reverb' really did. The fact of the matter is I was sort of half wrong and half right (depending on which end of the glass you call full). The DRE-S777 does use sampled reverb recordings, but they're not 'samples' in the same way that a sampler has samples – they're not fixed recordings that are played back at the push of a button.

In fact, Sony's use of the word 'sampling' is more related to when a chemist takes a sample of a substance for further analysis. Because that's how the reverbs in the DRE-S777 have been created. The ambience of a room has been recorded as it responds to different forms of audio signals, which provides the reverb 'sample' data. This data is then analysed to create a 'construct' of how that room responds to differing types of audio signals. This 'construct' is then stored in the DRE-S777, allowing it to apply the sampled room's ambience behaviour to whatever signal is passed through it.

A processor that records and analyses a room's ambience and then allows you to adapt that ambience to your own dry signals sounds like a terrific idea, but surely the end result will be too reliant on the quality of the original reverb recording? To solve that problem, Sony assembled a crack team of engineers from Sony R&D (and affiliated researchers) and sent them on a tour of great churches and concert halls known for their acoustic ambiences (stop-overs included Amsterdam, Vienna, Berlin and Spain). The reverb qualities of each church were sampled and analysed, and these are the 'spaces' found in the DRE-S777.

The Theory

The method Sony use to apply these reverb recordings to an input signal is called 'convolution'. We've all (I hope) walked into a space and clapped our hands to get some idea of the room's reverb characteristics. You clap and then hear the reflections of that clap bounce between the walls, floor and ceiling until the energy of that initial handclap dies. This clap sound can be considered as an impulse which contains a wide band of fre-

quencies, so it's a good start towards hearing how differing frequencies will behave in an environment. Sony have taken a 2kW PA system into their chosen halls, played computer generated signals through the system, and recorded the corresponding reverb. Back at Sony HQ, the reverb recordings are analysed to define how that room behaves under different signal conditions. This information is then loaded into the DRE-S777 and applied to whatever signals pass through it.

For example, let's assume a sampling rate of 44.1k. Think about how a single 44.1k sample word might sound by itself – very much like a handclap that only lasts 1/44,100th of a second. The DRE-S777 treats every one of those sample words as an individual handclap, and generates a reverb tail for each one based on the sampled reverb data taken from Sony's tour of famous churches and concert halls. So at 44.1k the device is computing a complete reverb tail 44,100 times each second (one for each sample word), and taking into account how the previous computations will affect those about to follow. As you can imagine, it requires a lot of processing.

It's a clever concept that requires the use of custom built Sony processing chips working with the data derived from the initial recordings. This is what gives the DRE-S777 the title of a 'sampling reverb', which could be slightly misleading. It really uses data collected from very good natural reverbs to compute reverb envelopes for every sample word it processes. You'll see from the following interview between Michael Gissing and Sony's Andrew Hingley that the potential exists for end users to sample their own reverbs for use in the DRE-S777, but most likely not for a few software revisions yet.

The Outside

When I first saw the DRE-S777 I was quite chuffed to see the imitation wood front panel. It gives the unit a look to suit its price tag, very much in keeping with where the DRE-S777 is aimed to compete, with the likes of the Lexicon 300L.

Michael Gissing talks with Sony's Andrew Hingley about the DRE-S777…

Michael Gissing: How would you describe the difference between synthesis reverb and the sampling reverb ?

Andrew Hingley: *The way to explain the difference is to relate it back to what happened to keyboards in the mid '80s. Original synthesisers came out 10 years earlier using subtractive synthesis, where you start off with basic waveforms and you filtered out frequencies, but then, with the Fairlight and E-mu Emulator, we saw the sampling keyboards arrive. So the easiest way to describe the difference with the DRE-S777 is to say that we are offering the same kind of change in technology as the change from synthesisers to sampling keyboards.*

MG: Except that keyboard sampling is the recording of a real sound in a digitised form, which is then played back from a keyboard command. But when we talk about a sampling reverb, you are accurately capturing the characteristics of an acoustic and using that to generate a reverb signal.

AH: *Well, it's based on a process called 'convolution'. The classic example is that, when you walk into a room and do a hand clap, you hear the reverb of the room responding to it. If you saw the envelope of that hand clap on an oscilloscope, or on a graph plotting time along the horizontal axis and amplitude along the vertical axis, you'd get a shape of the room's response to the hand clap. For the DRE-S777, we sample the room's response and feed that into our convolution 'engine'. Sampling theory says that if we then perform this convolution on an incoming signal, we will be effectively reproducing that reverb.*

If you imagine your digital audio samples as a series of spikes in the time domain, and each spike is the amplitude of each sample (represented by the numerical value of the digital audio signal), what you do is take the first sample into the convolution engine and create a whole series of samples which are the envelope of the room's response to the hand clap. Then you go forward one time period and multiply the second sample by that amplitude. The first sample has already generated a whole lot of values and then you add the appropriate values together which effectively is recreating the reverberant sound.

The difference between sampling reverbs and the more traditional synthesis reverb is very similar to the difference between synthesiser keyboards and sampling keyboards. The actual sampling process is different, but the quality difference between trying to mimic a sound using a synthesiser and actually sampling it is fundamentally on the same level. So it is accurate to say that we are sampling the room.

MG: So the only thing that sampling reverb does not accurately recreate would be the resonance of say a string instrument, played in a concert hall like the Concertgebouw, where the reverb causes the instrument to self resonate.

AH: *Well the analogy of sampling keyboards continues. If you listen to a piano sample, it sounds wonderfully like a piano, but it doesn't sound exactly like a piano. Our sampling system uses a PA and microphones with their inherent distortions. We choose where to put the speakers and the microphones so we are not trying to say that we are totally copying the reverb. We have inherent distortion in the measurement system we use, but the net result is an extremely natural sounding reverb. It sounds fundamentally different from synthesiser reverbs. But, in the same way that the piano sample depends on the type of microphone and its position, the sampling reverb's sound depends on the positioning of the PA and microphones within the room. Place them in different positions and you will get a different sound…*

MG: Yes, of course. I had the luxury of recording the Australian Youth Orchestra in the Amsterdam Concertgebouw, but my microphones and the mic placement were very different to the house microphone array. If they had lowered their mics into the recording position we would have had a very different recording.

AH: *Exactly right. Also, I think that for classical recording there is an important role for our reverb in post production. When editing 'takes' together or favouring a close mic, it would be useful to use our reverb with a sample of that room to add reverb that is very closely akin to the recorded reverb.*

MG: In a lot of mix environments the preference is to send an

auxiliary signal to the reverb unit and then mix back the wet return with the direct signal in the desk. Now given the huge processing that goes into convolution is there any significant signal delay?

AH: *The DRE is like all digital reverbs. They all have a processing time. And they all have A/D and D/A converters…*

MG: Not in my studio!

AH: *Well you are one of a small band. The DRE works in the frequency domain, not the time domain. Again this goes back to signal theory and our 'hand clap' and the reverb envelope. That description works in the time domain. But we actually work in the frequency domain because it reduces the amount of processing needed. You can do things called Fourier Transforms to move between the frequency domain and time domain. You can fully represent an audio signal either as a time sequence of samples, or as a number of frequency spectrums of the signal. Now there is some processing time required to convert the time domain signal to frequency spectrums and put them through our convolution engine, but listening to it doesn't seem to create any problems. It seems the design of the chip means that the processing time is fast enough to not matter.*

MG: There's a mixer and EQ stage in the DRE-S777. Is that EQ only applied to the reverb component?

AH: *Yes. One thing we have found is that when you hear the real reverb compared to synthesis reverbs, people naturally want to add a bit of tops to the reverb returns, but the real reverbs produced by the DRE are actually much more broad band. You do get a situation where the low end can get a bit wobbly, with standing waves bouncing off the architecture and so on. We sampled a church in Spain that was built in the eleventh century, with lots of pillars around and a lot of low frequency oscillation in the reverb. Now whether that is useful or not is for the producer to decide.*

MG: So if it is not wanted, you simply roll off the bottom end of the reverb return.

AH: *That's right. The other thing is that the analogue performance of the DRE is also very high end. You can switch it to a mode where it is an A/D and D/A converter. So for a project studio it might seem a bit expensive as a reverb, but when you consider that you can also use it as a high quality stand-alone A/D and D/A converter, it becomes more affordable.*

MG: I am sure a lot of audio professionals will want to know if they can record their own samples. The Holy Grail for film sound people is to be able to match post sync sound to location sound. Unlike the music world, most of the material we deal with is recorded in a diverse range of acoustic spaces.

AH: *The DRE-S777 has only recently been made available and is running version 1.0 software. There is a plan to allow the device to do its own sampling in the future. So anyone who owns one will be able to take it to venues and, by connecting the analogue outputs to speakers and feeding the analogue inputs from a microphone, record their own samples. But it does need a lot of care when it comes to collecting samples. If the sample is no good then the box won't sound good. There is a big opportunity for film and TV people to use this reverb for exactly the application you mentioned.*

And there is no reason why it is limited to just sampling rooms. If someone has an interesting spring reverb or plate reverb, they can sample that, too. Anything that resonates, like, say a pipe, can be sampled. It seems that we have a device that is capable of reproducing things so accurately, it isn't just room reverbs that we can emulate.

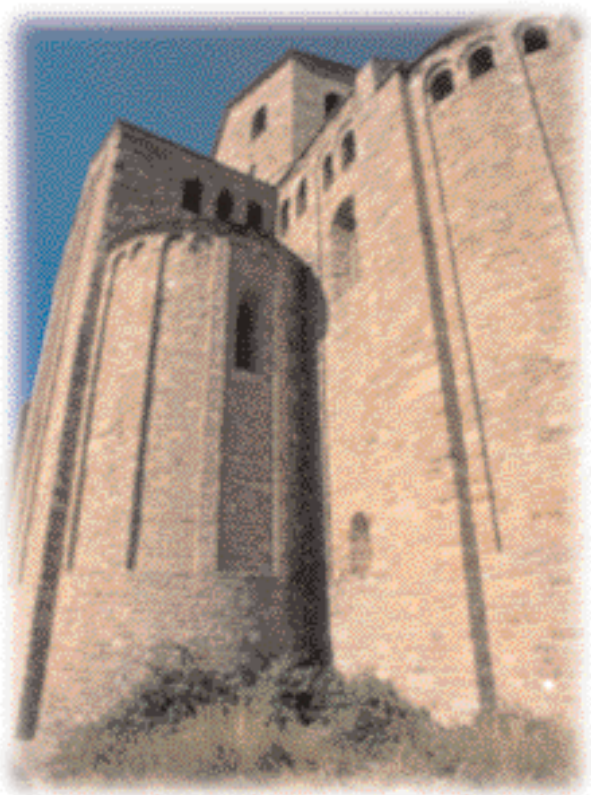

St. Vincente church in Cardona, Spain, one of the medieval acoustic spaces sampled by the Sony team.

MG: Like the sound of a wire fence?

AH: *…or underwater. The test tones that we use are actually swept tones, so anything you can put transducers on or in to generate and pick up that swept tone could end up as samples in the DRE-S777.*

And, of course, we can sample any device that has analogue inputs and outputs, like synthesised reverbs or flangers/phasers, although there are some inaccuracies in the sampling process and you are probably better of using the original device.

But although the sampling process introduces some inaccuracies, and there is not much to be gained by sampling synthesiser devices, there is one thing about the sampling process that makes it better. When we take the samples, we actually take lots and lots of them and run the tests many times, like a time lapse photograph, and if you do that you can average out the samples. The actual response you want always adds constructively, but background noise is randomised. So by taking lots of samples we can average them out and get a better signal to noise ratio. We sampled a plate reverb and, by using about 16 samples, we got 10dB better signal to noise ratio than the original device. So there is some validity in sampling old tape or spring reverbs. But for sampling modern digital reverbs or flangers, the only advantage you would have is the convenience of having it all in one device.

The DRE-S777 is two rack units high and has a very simply laid out front panel. On the left is a CD drive for loading the reams of data required. Above that is a PCMCIA card slot that, with the help of an ancillary card, accepts Sony's new 'Memory Stick' data storage technology. The Memory Stick is quite cute – up to 64MB of data in a card no larger than your house key.

The right hand side of the front panel is home to a large backlit LCD display with four software keys beneath it. To the right of the display is a very large data wheel that wouldn't look out of place on the dashboard of your Daimler Super V8. Overall there's not a lot to look after when driving the DRE-S777, definitely the mark of a professional tool. The unit is very deep, over 500mm, and weighs about 15kg with all the optional cards installed. Yes it's heavy, but so is a Daimler!

The Insides

You get some idea of the amount of data the DRE-S777 has to churn through after you've booted it up and found yourself waiting a good five or six minutes for all the reverb information to be loaded from the CD-ROM! This fills seven of the nine memory locations, and you get two medium halls, two churches, a 'studio' 'verb and two plates. Once you have the settings loaded from the CD-ROM it doesn't take long to find your way around, and there really isn't a lot to take care of.

Reverb time, or decay, is adjustable between 0.3 seconds and the length of the reverb you are using (determined by the room you have chosen). The only other parameter for the actual reverb is pre-delay.

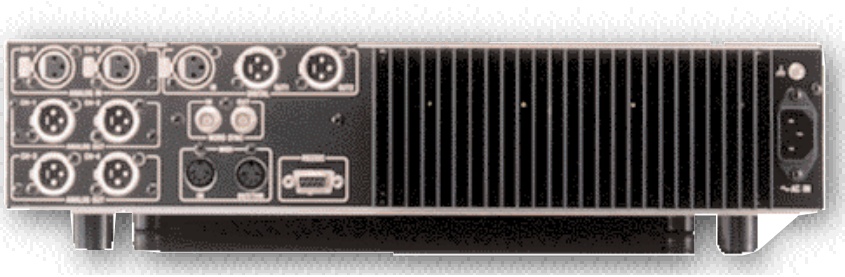

The Sony DRE-S777 back panel, packing a complete set of optional I/O.

After the reverb settings is a mixer page where wet and dry signals are adjustable and separately mutable. Following the mixer is an EQ page. This is again relatively simple with a high and low shelf and two parametric mid bands. EQ can be applied to either the wet or dry signal, or applied to whatever arrives at the input of the machine.

There's not a great deal to look after parameter-wise in the DRE-S777. The spaces sound so real and spacious that these few simple parameters are all you'd want. Come to think of it, I have three or four reverbs that I use with the machines in my studio and that's about all I adjust anyway: pre-delay, reverb time, and then EQ the result back at the console. It's nice not to be burdened with diffusion, density, spin, or width functions. The DRE-S777 sits there merrily convoluting its spaces to your program material with a delightful absence of finicky options.

The Ins and Outs

With a box of this calibre you'd expect to see some top notch analogue to digital conversion occurring. The DRE-S777 is certainly not a disappointment, and I'd have no doubts about using the A/D converters for high quality analogue conversion. Signal to noise ratio and dynamic range are quoted as 110dB. In its standard configuration the DRE-S777 will run as a mono in/stereo out or full stereo device at 44.1k and 48k sampling rates. One XLR AES/EBU digital input and two outputs also grace the standard configuration. These will accept 24-bit data, but using 88.2k or 96k sampling rates requires the addition of optional cards. I don't think making 96k an optional feature will upset anyone. Most facilities are only just reaching 96k capability, and any prospective buyer of the DRE-S777 would appreciate the choice and lower cost of the 44.1/48k unit rather than pay for the unrequired 96k spec.

And the sound?

That's the boring stuff out of the way. How this reverb sounds can only be described as exceedingly real. Lexicon really do have something to worry about here. The depth of the image is quite amazing. Reverb tails don't so much as dissipate as melt into nothingness, just as real world reverb does. It's remarkably easy, with a light touch of EQ, to pull some quite staggering spaces.

This really is a new angle for reverb processing and I think the method will be as revolutionary as the 16-bit sampler uprising of a few years ago. I won't be rushing down to the bank increasing my credit to buy one because it's just a tad outta my league. But it has been awfully splendid having one of these rare (at the moment) machines hanging around the studio. For an in-depth look at where the DRE-S777 will find its vocation, check out the following interview between Michael Gissing and Sony's Andrew Hingley. In the meantime I'm getting my microphone technique ready for what could be a very worthwhile direction in effect and ambience processing.

Distributed by

- *Sony Australia Pty Ltd*

Phone: +61 (0)2 9887 6666
Fax: +61 (0)2 9887 4605

Price

- *$8600 (DRE-S777);*

$1600 (DABK-S701 A/D converter board);
$1100 (DABK-S702 D/A converter board);
$3100 (DABK-S703 Expansion DSP board, stereo input, 4-channel output, 88.2k or 96k Fs);
$1150 (sampling reverb software).