1985 — 2005 Twenty years of the world's best music recording magazine

SOUND ON SOUND

The World's Best Music Recording Magazine

USA/International Edition

AWS900

An SSL for the home studio!

The biggest & best for reviews

Ableton *Live 5*
Korg OASYS
Digidesign M Box 2
and many more...

Hendrix in the studio

Eddie Kramer on this month's Classic Track

DVD ROM WITH THIS ISSUE
DVD002
www.soundonsound.com

Lexicon MX200 The hardware plug-in reverb ■ Waves *GTR* Amp modelling software ■ The Go! Team ■ Patchmix masterclass Get the best from Emu's soundcards
Crane Song Avocet Super monitor controller ■ *Cubase*'s video features explained
■ Make any hardware controller work with *Logic* ■ Kawai MP8 Stage piano

11
0 74470 05361 3
USA $9.95 New Zealand $14.99
Canada C$12.95 Australia $13.95

MIDI Velocity Mapping In Software Samplers

Reprint from Nov. 2005 issue of Sound On Sound. (c) Sound On Sound
(DVD article refers to audio examples found in the accompanying Nov. 2005 DVD)

Do all software samplers accurately map MIDI velocity? Is EXS24 better or worse than Kontakt or GigaSampler? Our guide reveals the dynamic truth. Play the embedded audio examples, as we explore one of the less documented aspects of using software samplers triggered by MIDI.

By Ernest Cholakis

When attempting to create a realistic MIDI performance using drum samples, I discovered huge discrepancies between the MIDI performance and the actual original acoustic performance the MIDI file was based on. To get to the bottom of this, I decided to graph the dynamic range of real-world instruments versus the dynamic range of samplers. This allowed me to visualise the nature of the transfer function (the relationship between the output and input signal). I was quite surprised by how varied the results actually were.

To illustrate the results, I have provided graphs and audio examples to document the comparative differences between the acoustic performances of electric guitar, drums and piano, and the same 'digital performances' triggered from the sampler. From this research, I've explored the use of new technical features and the creative use of existing features that might be incorporated into current samplers to correct the existing limitations that are discussed in this article.

Dynamic Range Of Acoustic Instruments

So what really is the dynamic range of an actual acoustic instrument? To answer this question I decided to record a piano, an electric guitar, a snare drum and a bass guitar, and then convert these recordings into graphs that visually illustrate exactly what the ear perceives. It's worth noting that this process is quite different from the waveform display of a sound editor, which shows volume in a linear manner, which is not how the ear perceives volume. Instead, the ear hears volume using a logarithmic scale, and because of this, I converted the amplitude scale (the vertical or 'y' axis) of all of the waveforms to the VU meter scale.

First, a series of piano notes at A4 pitch was recorded by continuously playing from ppp (pianissimo), gradually hitting the note harder in an even manner until fff (fortississimo) was reached. In Fig.1 we see the result which is very close to a straight line representing a total dynamic range of 41 decibels (dB). Fig.2 (a D5 note) has a 46dB range and Fig.3 (a G2 note) has 35dB. Fig.4, with a 26dB range, was generated by a Yamaha C7 Disklavier triggered by a MIDI file playing note A5 at 16 ascending velocity values.

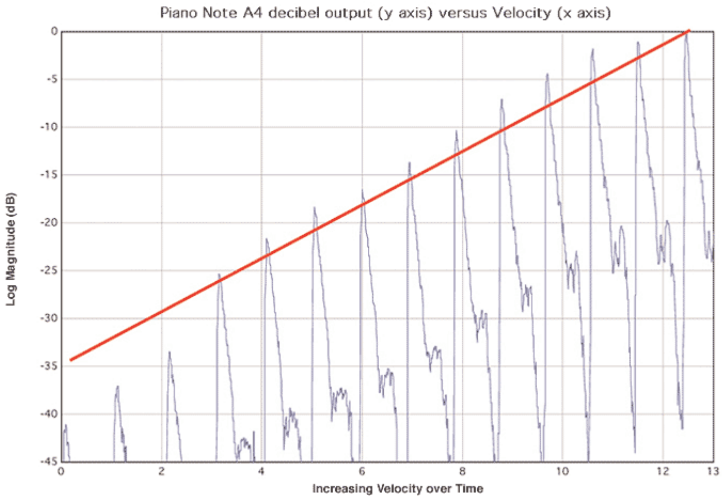

Fig.1: Decibel output (vertical 'y' axis) versus Velocity ('x' axis) of a series of acoustic piano notes (all at pitch A4) with ascending dynamics.

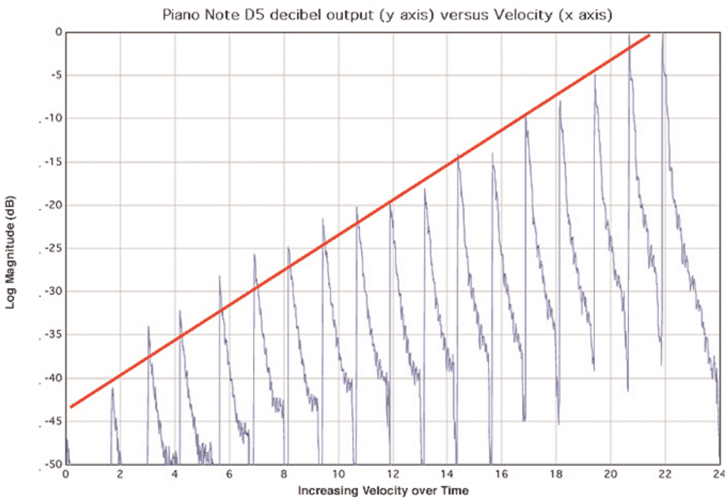

Fig.2: Decibel output versus Velocity of a series of acoustic piano notes at pitch D5.

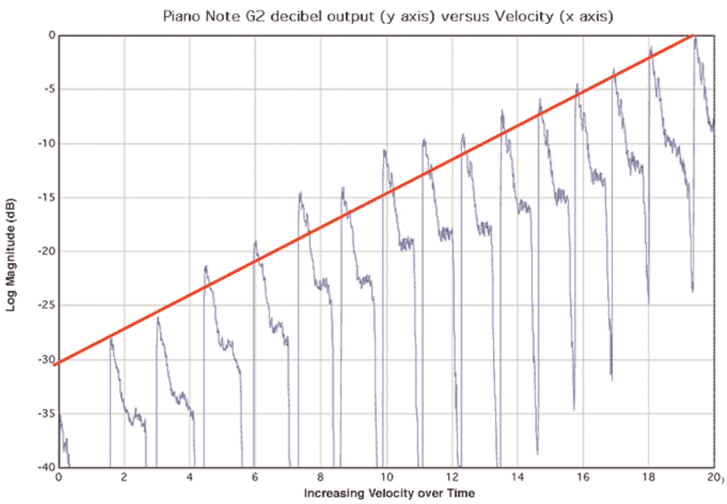

Fig.3: Decibel output versus Velocity of a series of acoustic piano notes at pitch G2.

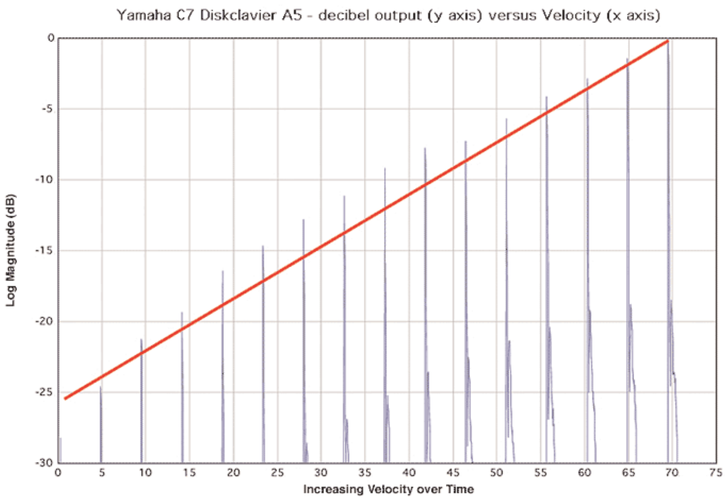

Fig.4: Decibel output versus Velocity of a series of Yamaha Disklavier notes at pitch A5.

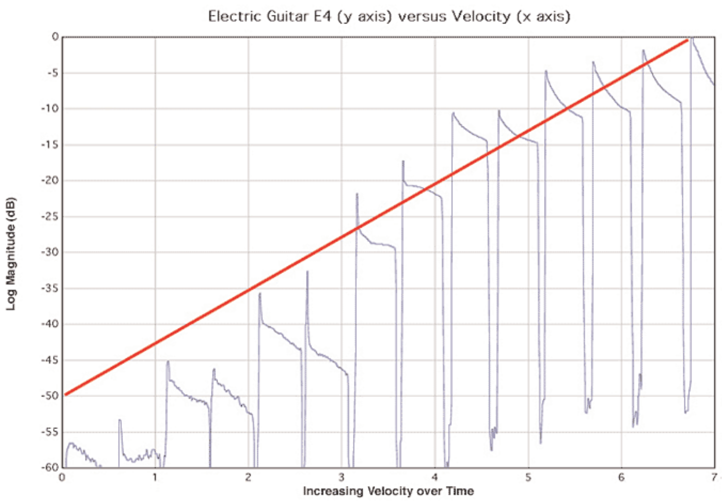

Fig.5: Decibel output versus Velocity of a series of electric guitar notes (all at pitch E4) with ascending dynamics.

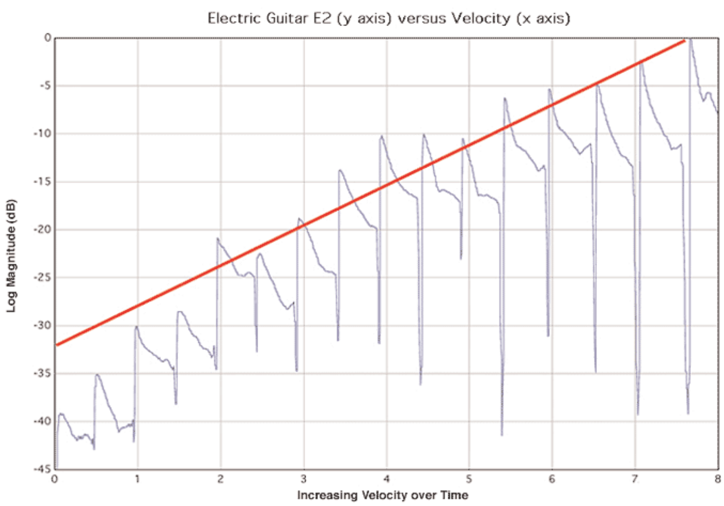

Fig.6: Electric Guitar notes at pitch E2.

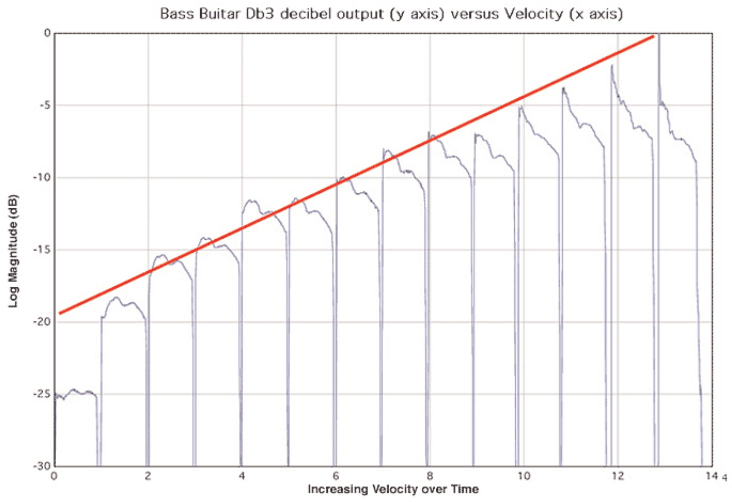

Fig.7: Decibel output versus Velocity of a series of Bass Guitar notes (all at pitch D3) with ascending dynamics.

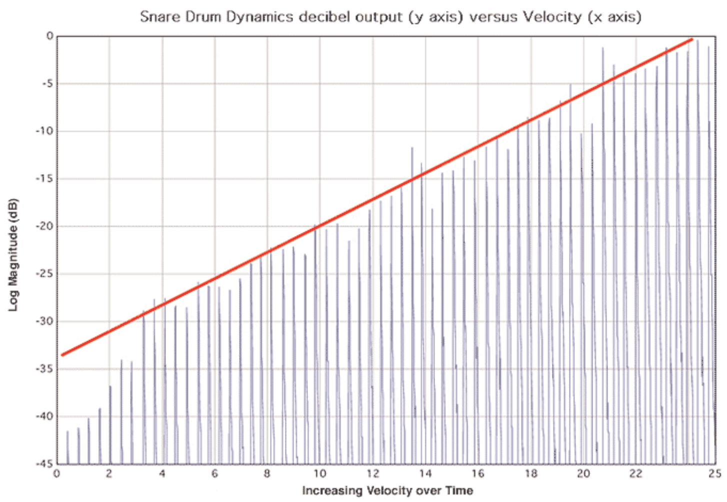

Fig.8: Decibel output versus Velocity of a series of Snare Drum hits with ascending dynamics.

When I recorded a plucked electric guitar playing the notes E4 (Fig.5) and E2 (Fig.6) a dynamic range of 56dB and 39dB, respectively, was attained. The bass guitar, at D3 (Fig.7), has a range of 25dB. A snare drum (Fig.8) had a dynamic range of 42dB. When one looks at the dynamic range in the graphs of each of these instruments, they all show a common linear pattern that is referred to as a 'linear decibel trend line'. In order to clearly illustrate this trend line I superimposed a 'best fit' red line onto each graph.

In making these recordings to illustrate the dynamic range, I realised that it was quite a challenge for a musician to play louder in very gradual increments, especially when attempting to play more than eight consecutive notes without having a single event's volume out of sequence. In the first example (Fig.1), 14 events covered a 41dB range with an average level increase of 2.9dB per event. It occurred to me that this is a good 'acid test' for determining a player's 'touch'. I suspect that a great player would be close to the trend line more often than not.

MIDI Velocity: Drum Analysis

Analysing A Real Drum Performance

I also thought it might be interesting to look at the dynamic range of the playing of a great drummer, such as Bernard Purdie (www.bernardpurdie.com). Fig.9 illustrates a four-bar groove (at 107 BPM) displayed in decibels. Typically, it's the snare drum that has the widest dynamic range, and in this example one can see several snare events peaking in the 0dB to -3dB range as well as the low-level snare hits in the -35dB range. Two examples (short spikes) of this can be seen between the 6.5 and 7 second range in Fig.9. Note there are several other snare hits in the -27dB to -34dB range. This low-level snare work is essential to the groove. Removing this or raising the level of such hits would ruin a great drummer's feel (as Figs.20-24 ably demonstrate later in the article).

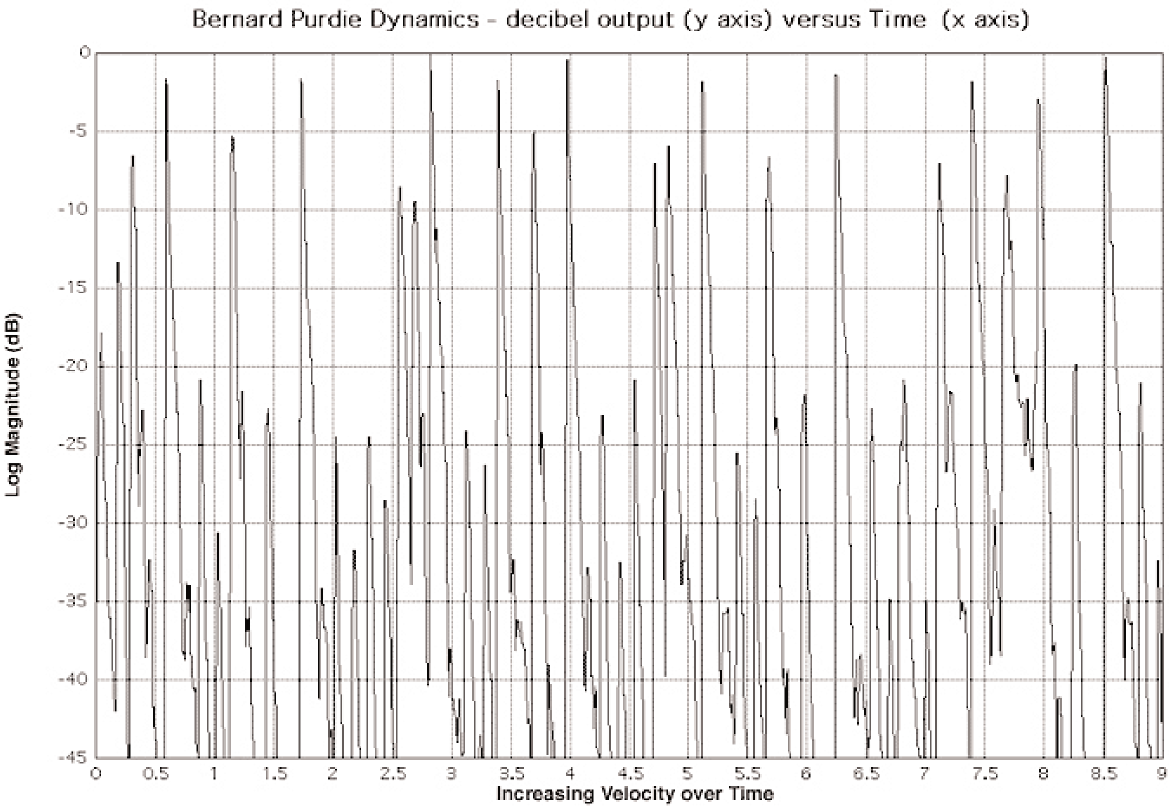

Fig.9: Graphic representation of the dynamics of a Bernard Purdie drum performance.

MIDI Velocity: Sampling Accuracy

Do Samplers Accurately Reflect The Acoustic Performance?

It is apparent from the earlier acoustic recordings that a linear decibel curve in the 40dB to 45dB range is needed to accurately reflect an acoustic performance. So, let's see how a selection of popular software samplers perform in this respect.

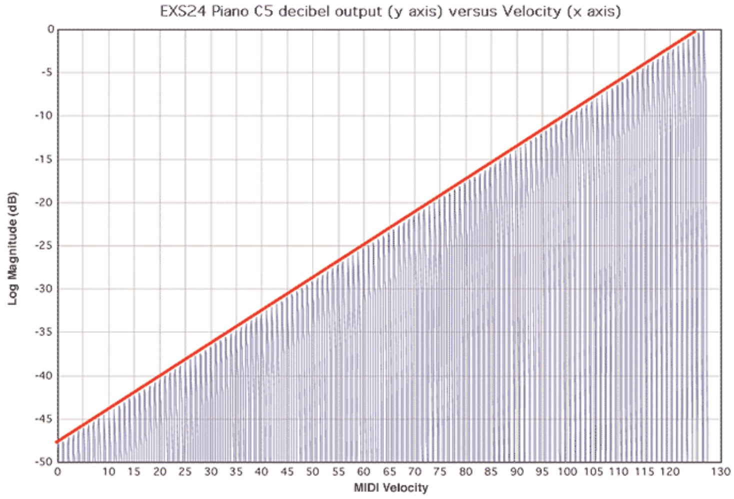

Fig.10: Apple Logic Pro EXS24 — The only software sampler with a true linear decibel curve.

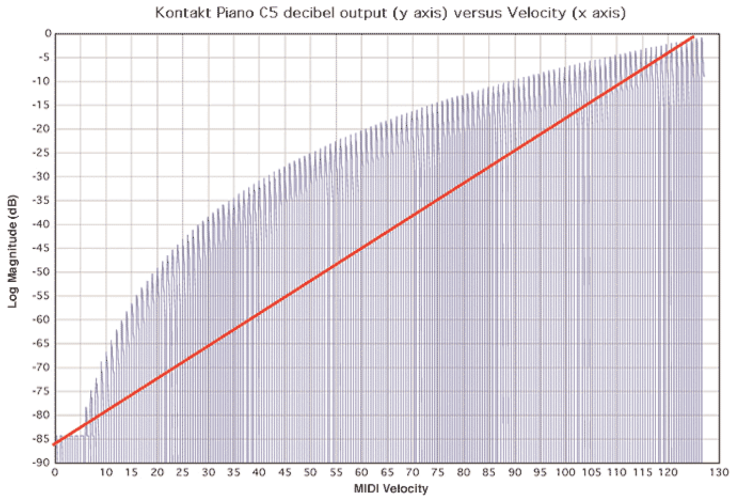

Fig.11: Native Instruments Kontakt.

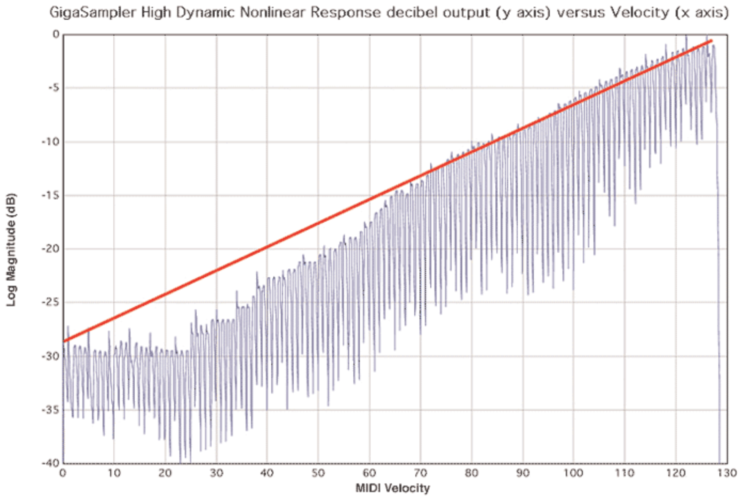

Fig.12: Tascam GigaStudio (in the non-linear high-dynamic range mode).

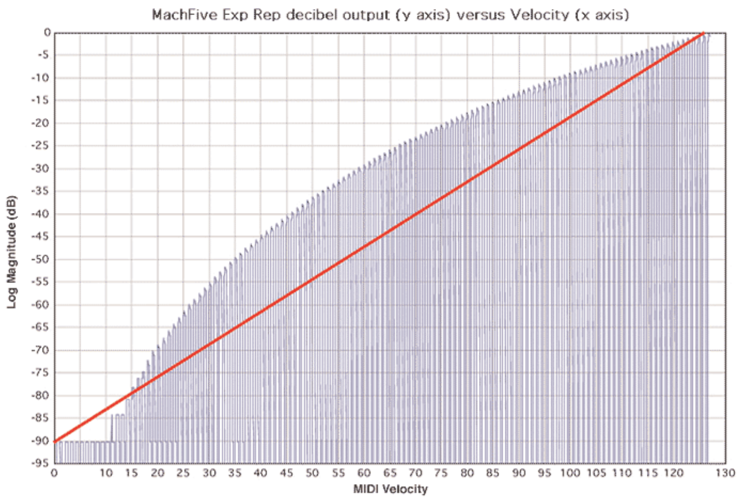

Fig.13: MOTU Mach Five.

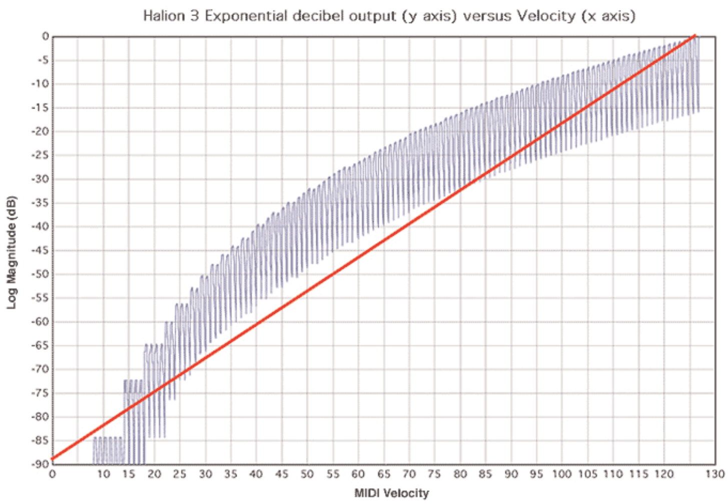

Fig.14: Steinberg Halion 3.

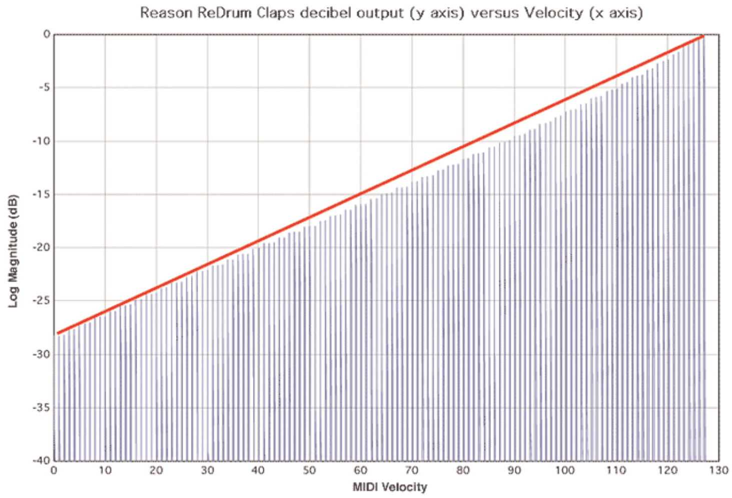

Fig.15: Propellerhead Reason ReDrum.

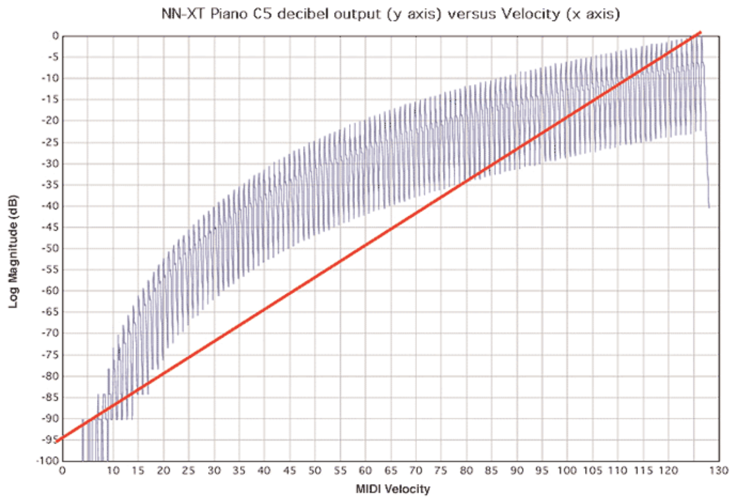

Fig.16: Propellerhead Reason NN-XT..

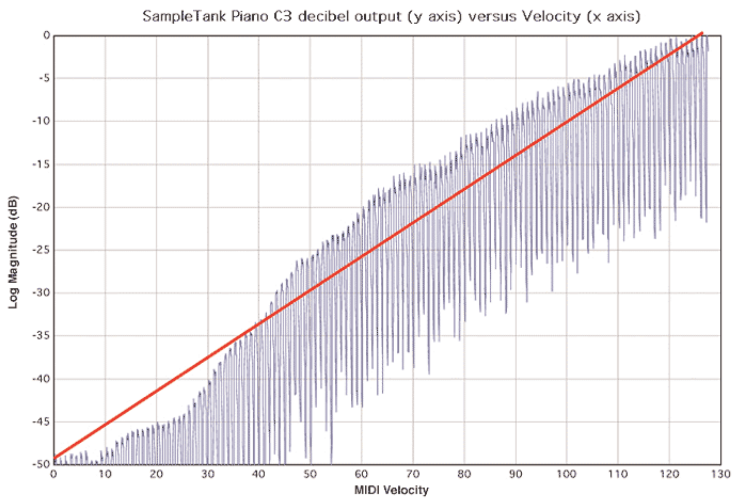

Fig.17: IK Multimedia SampleTank..

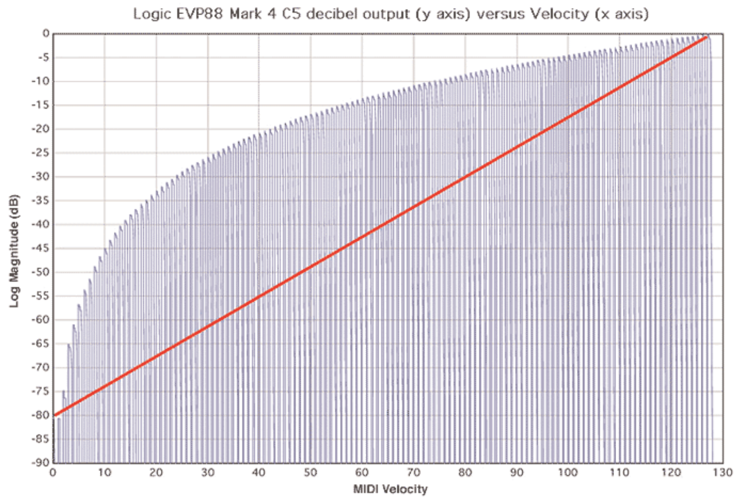

Fig.18: Apple Logic Pro EVP88.

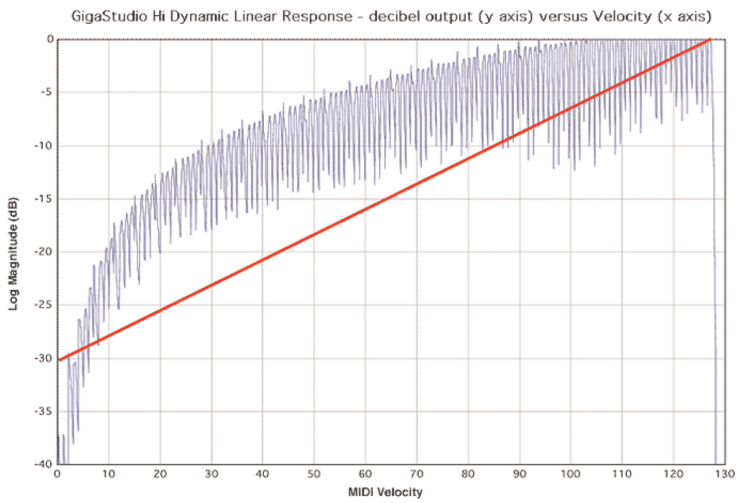

Fig.19: Tascam GigaStudio (in linear high-dynamic scaling).

The only software sampler with a true linear decibel curve is Apple Logic Pro's EXS24 (Fig.10). Additionally, in the EXS24 you can change the linear decibel curve over any range from 48dB up to 1dB.

Several samplers have a single fixed curve such as Propellerhead Reason's ReDrum, (Fig.15) which has a linear dB curve with a 28dB range. Tascam's GigaStudio (Fig.12) is close to a linear decibel curve when it is in the non-linear high-dynamic range mode, but its dynamic range is only 27dB. Clearly, Halion 3 (Fig.14), MOTU's Mach Five (Fig.13), Native Instruments' Kontakt (Fig.11), IK Multimedia's SampleTank (Fig.17), Propellerhead Reason NN-XT (Fig.16), Apple Logic Pro's EVP88 (Fig.18) and GigaStudio in linear high-dynamic scaling (Fig.19) are clearly not 'linear decibel' in their transfer functions.

	MIDI Velocity 10	MIDI Velocity 20	MIDI Velocity 40	MIDI Velocity 60	MIDI Velocity 80
Kontakt	-66dB	-49dB	-30dB	-20dB	-13dB
GigaSampler linear	-23dB	-16dB	-10dB	-7dB	-5dB
GigaSampler non-linear	-28dB	-30dB	-23dB	-18dB	-10dB
Mach Five exponential	-71dB	-52dB	-35dB	-22dB	-14dB
Reason NN-XT	-85dB	-52dB	-32dB	-20dB	-12dB
Reason ReDrum	-27dB	-24dB	-20dB	-16dB	-12dB
SampleTank	-49dB	-46dB	-33dB	-21dB	-12dB
Logic EVP88	-45dB	-34dB	-21dB	-14dB	-9dB
Halion 3	-84dB	-65dB	-40dB	-27dB	-16dB

Table 1. Apart from the highlighted white cell values, all the samplers produced different dynamic levels when receiving the exact same MIDI Velocity value. This should not be the case.

In Table 1 the actual output levels of MIDI velocities 10, 20, 40, 60 and 80 are displayed in decibels (for eight software samplers). Except for two matches at velocity 20 and 60 and three matches at velocity 80 (respectively highlighted in white below), all the dynamics levels are different. The individual dynamic range curves of the samplers make it virtually impossible to easily transfer a good MIDI performance created on one particular sampler to a different sampler. One could argue that this is needless technological complexity.

MIDI Velocity: Audio-to-MIDI Extraction

Audio-to-MIDI Extraction

One of the practical difficulties in extracting MIDI files from audio recordings is using a velocity value that will match the actual level of the acoustic performance. Without a linear decibel range of at least 40dB it is virtually impossible to achieve — that was until Logic Pro's EXS24 added this feature.

To illustrate this point, I extracted a MIDI file (Fig.36) from the Bernard Purdie four-bar drum loop that we heard in Audio Fig.9 and tried to recreate this performance with a different sounding drum kit (Audio Fig.20).

Bernard Purdie Drum Groove — at 40dB, 30dB, 20dB, 15dB and 10dB dynamic range. Kick Drum — at 40dB range HiHat — at 40dB range. Snare Drum — at 40dB, 30dB, 20dB, 15dB and 10dB dynamic range.

Next, individual drum parts are separated into Kick (Audio Fig.25), HiHat (Audio Fig.26) and Snare (Audio Fig.27). The snare track dynamics are then reduced from the original 40dB down to 10dB (Audio Figs. 27-31) as well as in the mix (with the HiHat and Kick) in Audio Figs. 20-24.

It is interesting to hear how the reduction in snare dynamics also reduces the smooth feel, by raising the volume of low-level events. The drum part begins to sound busy and start to leave less sonic space for the rest of the band. To hear this clearly, listen at a level that we hear when a real drummer is in the room — then try playing along! I suspect that you will find the 40dB snare track (Fig.27) leaves more room for your own performance. Another test that you might try is to play back the MIDI file (Fig.36) with your favourite sampler and see how it compares to the 40dB reference of Audio Fig.20.

MIDI Velocity: Dynamic Resolution

Dynamic Resolution: The Key To Realism In Sampled Performance

So what are the sequencing challenges when working with non-linear decibel curves (ie. all the samplers tested, with the exception of EXS24)?

From a musician's intuitive playing perspective, these curves give the wrong acoustic 'feel'. When playing either soft or loud passages on a keyboard, the actual levels heard do not match any real world acoustic instrument. The pianissimo (ppp) quiet passages are typically heard too loud and fortissimo (fff) passages aren't loud enough. Programming realistic low-level snare passages is often a challenge because there are not enough velocity values in the -30 to -35dB range — often only five values — and this doesn't give the user enough dynamic resolution to work with. There are workarounds, of course; one could split the snare parts into two separate tracks — one containing the high-level notes and the other the low-level notes — then adjust the low-level snare's volume at the mixing stage. But again, isn't this needless complexity?

A sampler with a variable linear decibel transfer function has several important advantages. One can compress a good performance by simply changing the decibel curve to a smaller value (as can be achieved easily in Apple Logic Pro's EXS24 sampler). Secondly, adding or subtracting a constant MIDI velocity amount to a sequence will not change the relative volume between each note in a MIDI file (as, of course, it does with non-linear decibel curves). Sequencing a drum accent at various levels requires a single fixed number when using linear decibel curves. With other curves, it depends on where you are working. For example, a 3dB accent may require moving the velocity level from 90 to 110, but if the accent is required for a note at a velocity of 30, it may work out to be 31 or 32, as shown earlier in Table 1 (see 'Sampling Accuracy') — and to further complicate matters, the amount is different for each sampler!

Another important advantage to linear decibel curves is that they enable one to adjust the dynamic range of a sequence in a very musical manner. For example, a performance that has a 40dB dynamic range (Audio Fig.27 below) can be changed into a performance with a 20dB range (Audio Fig.29 below) simply by changing a single sampler parameter (as can be achieved in EXS24). Even though the dynamic range is reduced, all the relative levels of each note are still preserved but with half the volume depth of the original. The audio track would probably not need a compressor/limiter to reduce the dynamic range in a MIDI-based performance and the overall benefit of this is a more natural-sounding track.

This is impossible to do with all the other types of dynamic range curves. For example, listen to the original extracted snare part as its dynamics are progressively reduced on an ESX24 from the 40dB of the original (Audio Fig.27), to 30dB and ultimately to 10dB (Audio Figs.28-31 above). Then compare it with a 2:1 and 3:1 compression with Peak and RMS level detection (Audio Figs.32-35 below). It is apparent that the linear decibel reduction of this sampler is more neutral and transparent than when one uses an audio compressor to do the same function.

Essential Features For Realistic Acoustic Emulation

Clearly, having variable linear decibel curves for a sampler or synthesizer is essential for realistic acoustic emulation. A future that offers users the ability to input individual decibel values for each note of the scale will tighten the gap between the sampler and acoustic realms. There are several possible ways to programme this feature into a patch; smaller values in the bass and greater values in the treble of the patch will result in a warmer sound. This range should have nuance and not a strict consistency in order to reflect real world acoustics. Notice in the first piano example the A4, D5 and G2 (Audio Figs.1-3 below) all have different values (41dB, 46dB, 35dB).

Another variation of this feature could be achieved by subtly adjusting the dynamics based on a particular scale or harmony. For example in the key of C, all the C, E, F and G notes could have a dynamic range slightly less than the other notes of the scale. This could possibly be used to emphasise the selected tonality because these notes could have a slightly stronger presence than other notes in this key. Such a feature, if added to software samplers, could open up the door to many interesting musical possibilities.

Audio Examples (Included on the DVD)

Listen to the audio of the Bernard Purdie drum groove discussed later in the article.
Audio Fig.1: Piano A4 notes (41dB dynamic range).
Audio Fig.2: Piano D5 notes (46dB dynamic range).
Audio Fig.3: Piano G2 notes with ascending dynamics (35dB dynamic range).
Audio Fig.4: Yamaha Disklavier A5 notes.
Audio Fig.5: Electric Guitar (pitch E4) notes (56dB dynamic range).
Audio Fig.6: Electric Guitar notes at pitch E2.
Audio Fig.7: Bass Guitar notes at pitch D3 with ascending dynamics.
Audio Fig.8: Snare Drum with ascending dynamics.
Listen to the audio of the Bernard Purdie drum groove illustrated in the Fig.9 graph below.
Fig.9: Bernard Purdie four-bar groove.
Fig.20: Recreation of Purdie drum groove using different drum kit (at 40dB).
Audio Fig.21: Drum groove recreated using samples, with dynamic range of the snare drum part reduced to 30dB.
Audio Fig.22: Drum groove recreated using samples, with dynamic range of the snare drum part reduced to 20dB.
Audio Fig.23: Drum groove recreated using samples, with dynamic range of the snare drum part reduced to 15dB.
Audio Fig.24: Drum groove recreated using samples, with dynamic range of the snare drum part reduced to 10dB.
Audio Fig.25: Kick drum part recreated with samples, using the 40dB dynamic range of the original performance.
Audio Fig.26: HiHat part recreated with samples, using the 40dB dynamic range of the original performance.
Audio Fig.27: Snare part recreated with samples, using the 40dB dynamic range of the original performance.
Audio Fig.28: Snare part recreated using samples, with dynamic range reduced to 30dB.
Audio Fig.29: Snare part recreated using samples, with dynamic range reduced to 20dB.
Audio Fig.30: Snare part recreated using samples, with dynamic range reduced to 15dB.
Audio Fig.31: Snare part recreated using samples, with dynamic range reduced to 10dB.
Fig.27: Snare part recreated with samples, using the 40dB dynamic range of the original performance.
Fig.28: Snare part recreated using samples, with dynamic range reduced to 30dB.
Fig.29: Snare part recreated using samples, with dynamic range reduced to 20dB.
Fig.30: Snare part recreated using samples, with dynamic range reduced to 15dB.
Fig.31: Snare part recreated using samples, with dynamic range reduced to 10dB.
Fig.32: Snare part compressed at 2:1 using Peak detection.
Fig.33: Snare part compressed at 2:1 using RMS (average) detection.
Fig.34: Snare part compressed at 3:1 using Peak detection.
Fig.35: Snare part compressed at 3:1 using RMS (average) detection.

Fig.1: Piano A4 notes (41dB dynamic range).
Fig.2: Piano D5 notes (46dB dynamic range).
\Fig.3: Piano G2 notes (35dB dynamic range).