DSD faq

Grimm AUDIO
www.grimmaudio.com
info@grimmaudio.com

Why buy a DSD converter? DSD is only for SACD, isn't it?

DSD is another word for 1-bit PCM sampled at 2.8224MHz. As a release format, SACD is the only one actually using this 1-bit data, but as an internal processing format it's surprisingly common. The vast majority of audio converters operate at 1 (or a few bits) at megahertz sampling rates. To get to PCM, digital low-pass filtering and decimation is required. This is mostly invisible to the user because converter ICs contain the decimation filter on-board, and put out PCM. But even if you're not aware of it, your PCM converter has quite a bit of DSD-like processing going on inside.

Alternatively the decimation can also be done externally. Doing so has the benefit of better control over conversion quality. The filters in converter IC's aren't optimised for sound quality but are designed for minimum latency and processing power. Several major DAWs sport DSD inputs, and their use is warranted even when doing a purely PCM production, simply because the decimation filters used by the DAW are vastly better than those found on-board PCM-output converter IC's.

All right, but why buy a 64fs converter? This seems like old technology. With 128fs and 256fs DSD and DXD (352.8kHz / 32-bit) formats being introduced, aren't these formats a safer bet?

The higher rate formats were designed to get around technical issues related to signal processing in DAW's, not to solve sonic issues in converters. The quality of the signal is ultimately limited by the analogue performance of the converter used, not by the data format. Grimm Audio strives to attain the best audio performance possible. The primary recipe is to keep the signal path simple, resulting in a discrete-circuit continuous time architecture. This design was shown to deliver its best at 1 bit and 2.8224MHz sampling rate. A substantially higher sampling rate would introduce settling time problems, resulting in noise modulation artefacts. More bits would introduce linearity problems and yet another breed of modulation artefacts. This strategy has paid off in unprecedented measured and sonic performance. Since the internal format of the AD1 is 1-bit 64fs data, it would make no sense to incur extra cost to provide interfaces with higher data rates.

DXD is a case in point. DXD is basically a 32-bit floating point format running at 352.8kHz. By converting incoming DSD to this format, all audio processing and storage can then be done using normal PCM processing without having to revert to 1-bit after each step. When Philips engineers were building DSD editing tools the design criterion was to find a PCM format that would not detract from the sonic capabilities of DSD. It was found that it was possible to convert a DSD signal to 352.8kHz/32 bit and back without incurring any audible quality loss, as long as good care was taken with the filtering and remodulation stages. It follows that using the same software it is also possible to convert from DXD to DSD and back without perceivable quality loss. If a converter or DAW cannot convert from DXD to DSD and vice versa transparently, this means the processing is incorrectly implemented. It can not be used to "prove" that DXD is more transparent sounding than DSD.

Granted that 1-bit conversion is the best choice for the A/D conversion stage, does it make much sense to store recordings in that format, considering it's most likely going to be converted to PCM at the editing stage anyway?

This depends on how your recording and editing sessions are structured. Certainly for recordings made without the immediate need for digital manipulation (like with classical music recording), the use of DSD as a storage format is warranted. Practice so far shows that DSD is at least as sonically transparent as 192kHz/24 bit and better than 96kHz/24bit. However, one channel of DSD takes up only 2.8Mbit/s, whereas one channel of 192kHz/24bit takes up 4.6Mbit/s. Given that the AD1 puts out DSD data anyway, it's most economical to store the audio in this format. Converting to PCM during recording would only increase the data rate without any added benefit. The conversion to PCM is best left to the DAW (or Grimm Audio's DD1) when the recording is loaded for editing.

Is the AD1 the optimum choice for all DSD-capable DAWs?

Yes. For a SADIE DAW it is the natural choice because the internal processing operates at 64fs. The AD1 operates at 64fs internally as well so it is the shortest possible link between an analogue input and a SADIE DAW. For those who prefer working on Pyramix and Sonoma systems it is also the optimum choice. Even though these DAWs offer 128fs DSD or DXD inputs in addition to 64fs DSD, there is no converter that is more sonically transparent than the AD1, regardless of whether outputs in the alternative formats are available from them. Users of non-DSD-capable DAWs will need to use an external decimation filter like the Grimm Audio DD1.

Is any of this interesting if I only produce PCM recordings?

By all means. It may be true that some PCM-only converters produce somewhat less noise above 20kHz, this design choice reflects a suboptimal compromise when actual audio performance is concerned, especially considering that most audio is still released as CD. There, the supposed benefit of lower HF noise is entirely lost while the performance loss inside the 20kHz range, incurred by optimising for out-band noise, is maintained.

When mastering an SACD, the final stage is a 1-bit noise shaper. I'm told it's not a good idea to run a signal through a noise shaper twice.

There are certainly problems (accumulation of HF noise) to be expected from applying several stages of deltasigma modulation to an audio signal. Twice is not a problem, five or six times is another matter. This means that using a deltasigma A/D in the front and a final digital remodulator at the back is not going to cause any trouble. This is especially so considering the particular choice of modulator used in the AD1. For one, as a continuous-time design it outputs slightly different data compared to a digital modulator, meaning you're not stacking two identical processes. For another, its outband noise characteristic is about the most gentle one found in DSD converters anywhere, largely owing to the lack of coefficient rounding. The HF noise floor of the digital remodulator is therefore always going to dominate the end result.

What is the max level of your DSD stream? DSD can peak above 0 dB I am told.

The absolute signal level of a DSD signal is expressed as modulation index. A 100% modulation index means all output samples are "1", -100% means all output samples are "0". When no signal is present, modulation index is 0% so on average half of the output samples are "1", the other half "0". Practical analogue 1-bit modulators produce negligible distortion when used below 50% peak modulation index. It is therefore common practice to design the decimation filters in 1-bit based PCM converters to output 0dBFS when the modulation index of the 1-bit modulator output is 50%.

This practice has become enshrined in the Scarlet Book standard, where the 0dB reference level for DSD signals has been set to 50% peak modulation index. The standard strongly recommends that the CD layer of hybrid discs be derived from the 2-channel SACD layer with 0dBFS = 0dBDSD. All SACD players are designed such that the signal levels in DSD and PCM modes match when this convention is followed. This does not mean, however, that the deltasigma modulator in DSD A/D converters falls apart above 50% modulation index, far from it. The AD1, for instance will provide acceptable performance for signal levels up to +1.8dBDSD. Above this level, modulator overload occurs. Using this margin to boost one's loudness level by 1.8dB is not a good idea though. In order to get the best commercial SNR spec, the D/A converters in most CD/DVD/SACD players max out at 0dBFS and hence 0dBDSD. Feeding them DSD signals modulated above 0dBDSD will cause analogue overload in the filters, digital overload in the preprocessing or both. Therefore, playing the Loudness Race game by modulating above 0dBDSD is liable to result in unplayable discs.

The level confusion is further compounded by the peak level spec set forth in the Scarlet Book. The Scarlet Book spec defines maximum allowable peak modulation in terms of a running flat average of DSD data bits. A pressing plant will refuse an SACD master if out of any consecutive 28 samples more than 24 or less than 4 are "1". In modulation index terms this corresponds to +/-71.4% or +3.10dBDSD. In keeping with the Scarlet Book standard, DSD DAWs sport a peak reading meter based on a 28-sample average.

This has led to the belief that an SACD can be made to play 3.1dB louder than a CD. This is quite incorrect, because a 28-sample unweighted average is a very shoddy low-pass filter and includes a very large amount of shaped HF noise from the modulation process. A typical 1-bit modulator will produce a "+3.1dB" reading on a 28-sample averaged level meter when the input level is around +1.5dB. By the same token, maximum modulation level specs given by several converter manufacturers are also misleading, because they state the 28-sample average reading at the 1-bit output, not the corresponding input level.

I am encountering problems when mixing tracks recording in DSD, there's hardly any HF headroom. Therefore I hesitate to record DSD streams.

There is quite some variation in the amount of HF noise put out by DSD A/D converters. Most converters use a front-end running at a different sample rate or word length and use a digital modulator to convert this to a signal compatible with the DSD spec. Quite often these digital modulators are designed for the lowest possible noise floor. Tested using digitally generated signals, better than -130dB in-band noise is common. This is quite suboptimal because only the signal-to-noise ratio of the analogue front-end needs to be accommodated, while the improved "virtual" noise level comes at the expense of greatly increased HF noise levels. A similar effect is encountered in the 1-bit remodulation stages in DSD DAWs.

Furthermore, not all DSD DAWs use the best processing filters. During the development of SACD and DSD, much effort went into designing filters with optimum cut-off and slope that were demonstrably transparent in listening tests while providing enough HF attenuation to deliver precisely the processing margin required by practical processing needs. If you find your DSD production system does not have the headroom needed for mixing large numbers of channels with substantial gain, this may be an indication of the DAW and/or the ADC not implementing the best possible filters and remodulation stages.

Bruno Putzeys and Eelco Grimm,
October 1st 2005.