



# Audio Engineering Society Convention Paper

Presented at the 110th Convention  
2001 May 12–15     Amsterdam, The Netherlands

*This convention paper has been reproduced from the author's advance manuscript, without editing, corrections, or consideration by the Review Board. The AES takes no responsibility for the contents. Additional papers may be obtained by sending request and remittance to Audio Engineering Society, 60 East 42nd Street, New York, New York 10165-2520, USA; also see [www.aes.org](http://www.aes.org). All rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.*

---

## Why Direct Stream Digital is the best choice as a digital audio format

Derk Reefman, Peter Nuijten

Philips Research Laboratories

Prof. Holstlaan 4, 5656 AA Eindhoven,  
the Netherlands

### ABSTRACT

In this paper, an overview of Direct Stream Digital (DSD) signal processing is given. It is shown that 1-bit DSD signals can be dithered properly, so the resulting dithered DSD stream does not contain audible artifacts in a band from 0-100 kHz. It is also shown that signal processing can be done best in a high rate, multi-bit domain. Arguments are given that the minimal frequency span needed to comply with the human auditory system is roughly 0-300 kHz. Following the signal processing, final conversion to DSD is made. It is demonstrated that Super Audio CD (SACD) is a very efficient *consumer* format: it is the format which, while maintaining all necessary psycho-acoustical characteristics such as high band width, filtering with wide transition bands *etc.*, uses the least bits from the disk; hence offering the longest playing time.

### INTRODUCTION

In the past few years, there has been an evolving trend in the audio world to move from the standard CD-format (*i.e.*, 16 bit resolution,

and a sampling frequency of 44.1 kHz) to other formats, which offer a higher resolution. A nice description of this otherwise rather vague definition is given by the AES high-resolution audio technical committee: every carrier offering more than 2 channels, each 44.1 or 48 kHz and resolution corresponding to 16 bit, is coined as high-resolution. Hence, high resolution embraces both multichannel recordings and audio formats which allow for a higher definition of the audio data *per sé*, such as SACD, which stores 1-bit words at a sampling frequency of 64

times 44.1 kHz, and DVD-A, which covers a wide variety of sampling rates (44.1/48 to 176/192 kHz) and word lengths (16-24 bit). Often, the SACD data format is called bitstream or Direct Stream Digital (DSD) to contrast it with the Pulse Code modulation (PCM) used in the CD and DVD-A format.

Though sigma delta modulation (or its predecessor, *Delta* modulation), the basic principle behind DSD, has been known for a long time [11], many of its intrinsic properties have not been well-covered. Though many models exist, they are flawed due to linearization the quantizer, or, if a linearization is not made, from the fact that only a first or second-order modulator is studied. This contrasts to the situation in the realm of PCM, where many of its properties have been studied in depth both experimentally and theoretically, for example, [18, 10, 6]. The purpose of this paper is to give an outline of the technical properties of the DSD format, and to present an overview of the signal processing required for preparing an SACD loaded with 8 channels of high resolution audio data.

### PROPERTIES OF THE BITSTREAM FORMAT

In this section, some properties of single-bit sigma delta modulation will be displayed. Though most of the results mentioned in this section are not new, we think they deserve mentioning because they are not well-known within the audio community as illustrated, *e.g.*, at the AES 2000 in Los Angeles.

#### Introduction to Sigma Delta modulation

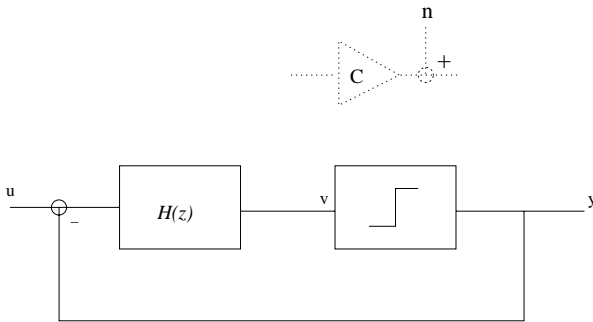


Fig. 1: Schematic representation an interpolative  $\Sigma\Delta$  modulator. Dotted is the replacement of the quantizer by an average gain  $c$  and addition of noise  $n$ .

In Fig. 1, the basic architecture of an interpolative  $\Sigma\Delta$  modulator is given. The quantizer  $Q$  is a two-level system; hence, the quantization error is very large, resulting in significant quantization noise. However, this noise can be *shaped* such, that virtually all noise power falls outside the range 0-20 kHz. For this reason  $\Sigma\Delta$  modulators are also called *noise shapers*<sup>1</sup>. The filter  $H(z)$  determines the character of the *noise shaping*; a good approximation to the noise transfer function ( $NTF$ ) of the modulator is given by

$$NTF(z) \approx \frac{1}{1 + cH(z)} \quad (1)$$

where  $c$  is the average gain of the quantizer. Hence,  $H(z)$  is usually a low-pass filter. Typically,  $NTF(z)$  is designed in such a way, that its -3

<sup>1</sup>In a strict sense,  $\Sigma\Delta$  modulators and noise shapers differ in topology; their characteristics, however, are virtually the same

<sup>2</sup>The reference 0 dB level is taken as the feedback level from the quantizer to the input. The standardized SACD 0 dB level is 6 dB lower.

dB cross-over point is around 70-90 kHz. The signal transfer is given by:

$$STF(z) \approx \frac{cH(z)}{1 + cH(z)} \quad (2)$$

A good introduction in the principles of interpolative (and other)  $\Sigma\Delta$  coders is given in [16]. In the sequel, we will focus on some properties of  $\Sigma\Delta$  modulators which are of crucial importance in audio applications.

One-bit  $\Sigma\Delta$  modulators (SDM's) differ from PCM AD convertors in the sense that the quantizer covers 2 levels only, and that the quantization error is fed back to the input. As a result, the output power of a SDM is always constant; its signal level is either +1 or -1, giving a signal power of 1. As known from PCM theory, dither is necessary to remove artifacts due the non-linear behaviour of such a device. The fact that the quantizer only spans two levels obviously removes the possibility to introduce the standard PCM dither [14]: dither with a triangularly shaped probability distribution function, with a width of 1 LSB, which clearly spans 3 levels. These two observations indicate that dithering a SDM is not identical to dithering a multi-bit quantizer. As dithering is a very important concept in linearizing both PCM and SDM quantizers, we will devote a section to this issue. In this paper, we will not address any theoretical issues, such as covered in [9], where it is shown that dithering higher than first-order sigma delta modulators is viable, but approach the subject from a more pragmatic point of view and show the results of various simulations. All simulations have been performed with a 7<sup>th</sup> order SDM, designed according to the general rules sketched in [1]. The design is based on a Butterworth design for  $NTF(z)$ , with a corner frequency of 70 kHz. By adding feedback over 2 integrators, a resonator section is created which allows to shape the noise in the baseband according to the general shape of the human hearing threshold curve. The topology of this SDM is shown in Fig. 2.

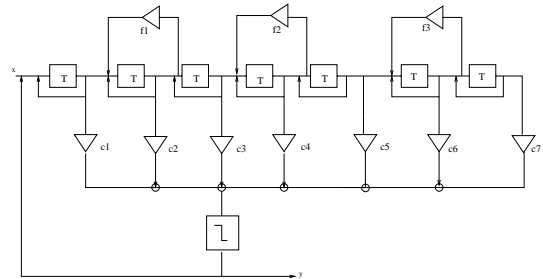


Fig. 2: A practical implementation of a 7<sup>th</sup> order SDM.

This SDM remains stable for inputs up to roughly -4 dB<sup>2</sup> and has a SNR of 126 dB in the band 0-20 kHz. Hence, it is a realistic SDM in the sense that it complies with the requirements set by the SACD specifications.

Note, however, that this modulator has not been optimized in any sense, unlike the modulators in use for SACD editing and signal processing which all have been fine-tuned for maximum performance. It is also important to realize, that all artifacts which occur in the baseband can be made almost arbitrarily small by increasing the order of the sigma delta modulator.

#### Dithering Sigma Delta modulators




Fig. 3: A practical implementation of dithering a SDM.




Fig. 4: In-band noise floor as a function of the input DC level. The lower curve corresponds to an undithered SDM; the upper to a dithered SDM.

A few aspects which bear a clear relationship to dithering are properties such as the in-band noise dependence on input signal level, harmonic distortion, and limit cycles ('birdies'). The general scheme that we use for dithering a SDM is depicted in Fig. 3. See also [15] and [16] for an overview of dithering SDM's.

The results that will be presented in the sequel, all refer to adding dither just before the quantizer (see Fig. 3). The dither has a uniform distribution, with an average amplitude approximately equal to the average amplitude of the input of the quantizer without dither.

In Fig. 4, the results for the in-band noise as a function of the DC input signal are depicted. For sine waves of non-zero frequency (0-85 kHz), the trends are virtually the same.

Clearly, the non-linearity of the SDM is demonstrated in the lower curve of Fig. 4. At inputs below -120 dB, the SDM is in a limit cycle, which is removed upon increasing the DC input to -118 dB. The

limit cycle causes high powered tones at 350, 700 and 1050 kHz. As before, since the limit cycle contained almost all the output power of the SDM, this power needs to be re-distributed as the limit cycle is broken up (the output power of a SDM is constant!). Hence, the noise floor in the signal band increases.

Interestingly, upon further increase of the signal, the noise floor remains approximately constant. Nonetheless, it is known that any SDM displays limit cycles for specific non-zero DC values as well, albeit that the power in these limit cycles is much less compared to the limit cycle of zero input. Inputs which are well-known to give rise to strong tones are simple rational inputs. This effect is illustrated in Fig. 5, for a DC value of  $1/250 = 0.004$ . In its simplest DSD representation, this would correspond to a sequence with a repetition period of 250 bits or 11.2 kHz, which is, indeed, almost equal to the result displayed in Fig. 5. A slight deviation from this simple reasoning occurs in high-order SDM's, as some overtones of the fundamental are visible.

The results in Fig. 4 for the dithered SDM show that, for exactly zero input, the limit cycle is absent, and the noise floor is at approximately -125 dB. On increasing the DC input value, the noise floor remains constant within 0.3 dB. Note, that the absolute level at which this variations occur can be made lower by increasing the order of the SDM; in fact, it is equally well possible to suppress the in-band noise to a level of -130 dB using a 7<sup>th</sup> order SDM. At this level, any variation by even 1 dB of the noise floor is rendered insignificant because it is completely overwhelmed by noise from the analog circuitry (not to mention the noise from microphones). Again, the effect of tones due to a small DC input value is completely removed by adding the appropriate amount of dither, as demonstrated by the upper curve in Fig. 5, at the cost of a slight loss in dynamic range. It must be noted, that the spectrum of the dithered SDM is, aside from some tones around  $f_s/2$ ,  $f_s/4$ , identical to the linearized response Eq. 2.

Even more interesting is the situation with respect to harmonic distortion. In Fig. 6, we have depicted the power spectrum of our undithered modulator.

Firstly, it should be remarked that the SDM depicted in Fig. 6 exhibits excessive harmonic distortion; most SDM's have a much lower distortion. The third harmonic is 12 dB above the noise floor; also the 5<sup>th</sup> harmonic is still very prominent. Nevertheless, if the  $\Sigma\Delta$  is dithered according to the scheme discussed previously, these harmonic distortion components disappear completely. The price to pay for the dithering is a slightly decreased dynamic range (the decrease is roughly 3 dB, from 126.5 to 123 dB), and a slightly lower maximum input: from -3.8 dB to -3.9 dB.

It is relevant to mention that these decreases in performance upon dithering are not specific to sigma delta modulation. Whenever a multi-bit quantizer is dithered to such an extent that its harmonic distortion in the signal band is absent (by, e.g., applying additive pdf dither [13]), its SNR decreases by 4.8 dB. Also, the maximum input range decreases slightly.

The dithering applied does not remove the spurious peaks around  $f_s/2$ , while it does so for much lower frequencies. This is in line with the general thinking behind noise shaping; any artifacts due to quantization are pushed up in frequency where they have little effect. A hand-waving argument that can be used here is that because of the large low-frequency gain of the loop filter, many dither samples are used to whiten the noise in the low-frequency part. In the high-frequency part, however, only few dither samples are used, which are not enough to completely remove artifacts from non-linear behaviour of the quantizer. If the analog system cannot linearly cope with the high frequency spectrum, these non-linearities can easily be filtered out digitally. This




Fig. 5: The effect of dithering on the in-band tones ('birdies') as caused by a rational DC input.

can be done without the need for steep filters (filters with a transition band from 100kHz to roughly 500 kHz are needed) to arbitrary levels, hence, we need not to worry about these.

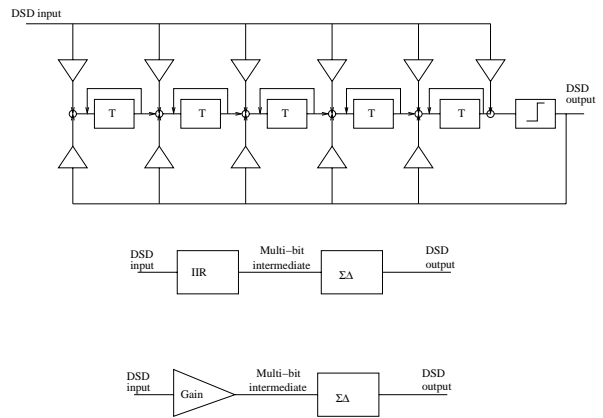


Fig. 7: Top: IIR filtering unit as elaborated on in [2]. Middle: the equivalent schematic, explicitly including the multi-bit path between the IIR filter and the SDM. Bottom: a simple signal processing step: gain adjustment, which is used as an example in the text

**SIGNAL PROCESSING FOR SACD**

An issue which needs to be addressed is signal processing for DSD streams. A number of papers have been published [8, 2, 3], which address signal processing in the single bit domain. Though the topologies shown in these papers may not directly reflect this, the general way of operation of these signal processing units is not different from the schematic in Fig. 7, which consists of an IIR filtering step, connected to a SDM via a multi-bit path. In DSD jargon, the SDM which is used to convert a multi-bit signal back to a single bit signal, is called a re-quantizer or re-modulator.

In actual applications, the signal processing will certainly not consist of only a single processing step; many steps will be cascaded to a sig-

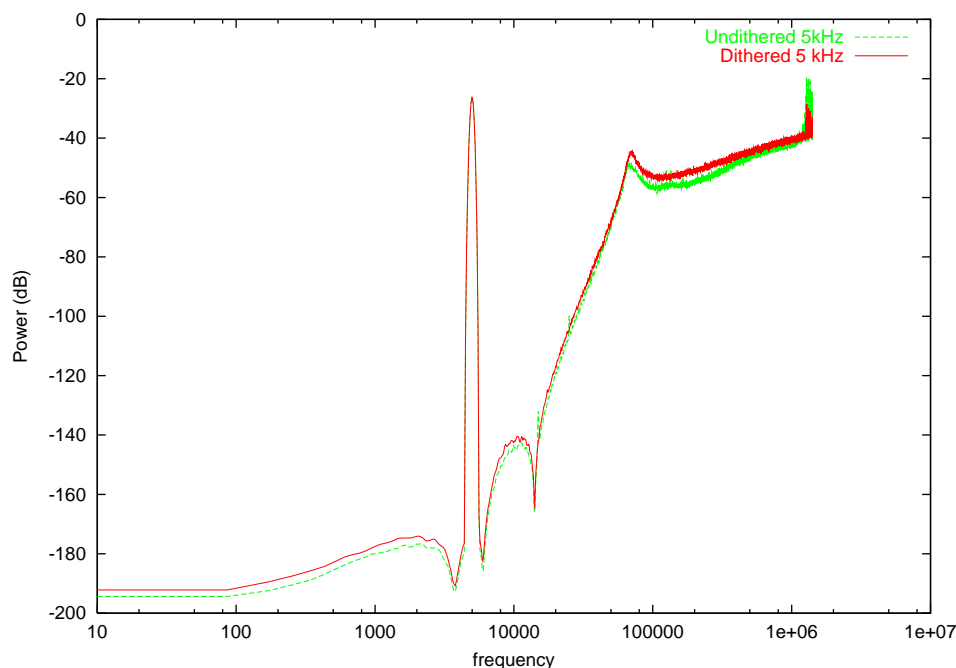


Fig. 6: Power spectrum of an undithered SDM, and a dithered SDM.

nal processing chain. It is this cascading which has its bearings on the viability of this approach.

For the purpose of illustration, we will explain the problems of multiple requantization using gain adaption (shown in the lower part of Fig. 7) as the signal processing step. For a standard (order  $> 2$ )  $\Sigma\Delta$  modulator, the maximum input level must be below 0dB; a reasonable value is  $-3$  dB. Suppose, that the gain value is set to  $-20$  dB. As the DSD input to the gain controller is 0 dB always, the input to the requantizing SDM is  $-20$  dB; a level that the system can deal with without problems. Now suppose, that the gain level is set to  $+3$  dB. The input to the requantizing SDM is now also  $+3$  dB, which will clearly cause overload of the system. It follows that the DSD stream has to be low-pass filtered in order to reduce the high frequency content, and creating an input power for the requantizer which does not exceed  $-3$  dB. Hence, this operation has its clear impact on band width. When after this requantizer yet another requantizer is to follow, the same operation has to be repeated, be it that the total power of high frequency noise has increased.

Hence, to keep the next requantizer in the stable operating regime, either the low-pass filtering applied has to be steeper, or the corner frequency has to be lowered. Though for a limited number of requantizations (say, 10) these band limiting operations have no serious effect - the bandwidth of 100 kHz is maintained, and the filtering has still a transition band of 100 kHz - for a larger number of requantizations it does have undesired effects. In current DSD signal processing approaches, the idea is therefore to limit the number of requantizations as much as possible.

This idea is used in, for example, the Merging Technologies[20] and the SONY Sonoma [21] editing systems. In these systems, the DSD signal is converted to a multi-bit signal, and all subsequent signal processing is done in that domain; only at the final stage, when all signal processing is done (after mastering) the signal is converted back to DSD.

As a consequence, the idea is to use either a 256 times 44.1 kHz single bit signal, which inherently has a high SNR to 90 kHz, or to use a multi-bit (say, 4) recording at high speed (64 times 44.1 kHz, equalling 2.8 MHz). All signal processing is subsequently done in the multi-bit domain, as demonstrated above, remaining at high sampling rates, and only in a final stage the conversion to 1-bit, 2.8 MHz is made. It is important to realize, that DSD for SACD is the final *consumer* format; it is the format which, while maintaining all necessary psycho-acoustical characteristics such as high band width, filtering with wide transition bands *etc.*, uses the least bits from the disk; hence offering the longest playing time. This brings us to the next section.

#### BANDWIDTH, DATA RATES AND DATA COMPRESSION

The large-input bandwidth of a  $\Sigma\Delta$  modulator is determined by the corner frequency of the high-pass filter, according to which the loop filter is designed. In this case this frequency is 70 kHz; hence for this particular SDM the maximum allowable full scale input ( $-4$  dB) is 65 kHz. For higher frequencies, input at lower levels (*e.g.*,  $-8$  dB) is still possible. This is illustrated in Fig. 8, where inputs with frequencies up to 130 kHz are displayed. Note, that due to the dropping loop gain, the transfer function of the SDM is straight up to 70 kHz, above which it slowly decreases. This explains the slightly lower amplitude of the high frequencies (75 kHz and higher) compared to the lower frequencies. The roll-off of this filtering, which is intrinsic to the SDM, is extremely slow. It is caused by the fact that the gain of the filter  $H(z)$  is dropping at high frequencies; hence, the signal transfer *cf.* Eq. (2) drops slightly. At 750 kHz the signal is only 10 dB down. Hence, it is clear that DSD as a carrier provides a signal bandwidth which exceeds 100 kHz by far. Obviously, the SNR decreases for high frequencies. As far as experimental evidence shows [17] the human auditory system is sensitive to ultra sound *only* if it is correlated with the signal in the 0-20 kHz band,




Fig. 8: Power spectrum of the SDM with inputs at  $-10$  dB, using a dithered SDM. Note, that the power spectra are independent of the various inputs.

hence, the increase in noise at high frequencies is not important. The excess in bandwidth can subsequently be used to allow for very slow low-pass filtering, which keeps the time response very accurate. Some more detailed explanations about the importance of bandwidth for the mere purpose of slow filtering is given in [7]. Indeed, filtering must still be applied to prevent tweeters in sound reproduction systems from overheating. This can be achieved easily with second-order filters, by which the level of high frequency noise is reduced by approximately 45 dB if the corner frequency is chosen as 100 kHz.

This effect is clearly illustrated in Fig. 9. The impulse responses of 4 different systems in a multichannel configuration are depicted: a 48 kHz system, with a bandwidth of 20 kHz (that is, 8 kHz transition bandwidth is allowed for anti-aliasing filtering), a 96 kHz system with 35 kHz bandwidth (26 kHz transition bandwidth), a 192 kHz system with 75 kHz bandwidth (42 kHz transition bandwidth) and an SACD system with 95 kHz bandwidth. Though none of the systems reproduces the input exactly, the DSD system shows the least artifacts. Clearly, the 48 kHz system has great difficulty in reproducing the click; due to the steep filtering it starts ringing at a  $-30$  dB level approximately 1 ms before the click, which is very audible. Also at the higher sampling frequencies, the ringing phenomenon cannot be removed, though it is reduced significantly. Only the DSD system is very effective in suppressing the ringing effect, due to very slow filtering above 95 kHz. The price to pay for this is the increase in noise floor with respect to the other systems; however, as the noise floor contains only high frequency components which are uncorrelated with the audio, they are not

perceptible.

Other issues which often appear to be confusing, are data rates in connection to the bandwidth claimed by SACD. The SACD format comprises (apart from its Red-Book conforming CD layer) two different music streams: a stereo 2 channel stream, and a surround 6 channel stream. Hence, an SACD contains 8 channels of high-quality audio. Because *all* channels are 2.8 MHz sample rate, 1-bit signals, the total data rate equals 2.8 Mbyte/s (or 22.4 Mbit/s).

On these signals, lossless coding is applied. This lossless coding scheme is specifically developed for coding 1-bit signals and is described in [5]. From experience of over 100 recordings, the average coding gain is roughly 2.4-2.5 for pop recordings, and 2.6-2.7 for classical recordings. This corresponds to a data rate per channel of about 1.1-1.2 Mbit/s. This indicates that on average 70 minutes of a DSD signal can be recorded on an SACD in the 8-channel format. For 6 channels, this amounts to roughly 95 minutes. Also, the high sampling rate of DSD allows for the use of filters with slow roll-off. We can compare this to DVD-A, much in the same sense as done in [12]. The DVD-A format that gets closest to the SACD characteristics is DVD-A at 192 kHz, 20 bit, which reaches the same dynamic range, but is either of lower bandwidth than SACD if sloppy anti-aliasing filters are used, or has the same bandwidth using steep filters. Using a compression factor of 2 the data rate amounts to 1.9 Mbit/s, which is almost twice as much as the data rate for DSD. Hence, even if only six channels are used on the optical disk (compared to 6+2 on SACD), only 55 minutes of music can be stored - much less than the 74 minutes that we are




Fig. 9: Responses (from left to right) of a DSD, a 192 kHz, a 96 kHz and a 48 kHz system on a -6 dB block input ('click') of 3  $\mu$ s duration.

accustomed to from CD.

#### AD AND DA CONVERSION AND OUTLOOK

The original idea for 1-bit coding stems from the observations in the 80-ties and early 90-ties, where the performance of AD convertors improved dramatically with respect to their multi-bit predecessors by using the  $\Sigma\Delta$  concept. Indeed, the DA conversion in the feedback path of a SDM is inherently linear, which caused a tremendous increase in overall linearity. Total harmonic distortion ratios of  $< -90$  dB became possible using this technique, concomittant with a dynamic range exceeding 100 dB.

With the increasing demand, however, for better performance, AD/DA manufacturers turned to multi-bit again. In this case not the 14-16 bit flash or folding architectures, but 4-8 bit noise shaping designs, which run at increasingly high speeds; 128  $f_s$  is no exception [19]. The sole reason for returning to multi-bit is in the fact that these designs are less sensitive to clock-jitter; the linearity problems have been effectively solved for these few-bit designs by advanced calibration techniques, or the use of dynamic element matching [4].

From this viewpoint, the best possible scenario is to use the native AD format itself as a pro-audio format: high speed, few bits. After all necessary signal processing, this format is converted to the DSD format, which maintains all necessary psycho-acoustical characteristics such as high bandwidth, filtering with wide transition bands *etc.*, while using the least bits from the disk.

#### CONCLUSIONS

It is shown that DSD signals can be produced that can be properly dithered. Here, properly dithered means that non-linear artifacts caused

by the 2-level quantizer can be removed from the band 0-100 kHz. Moreover, it is argued that the DSD format complies with the minimal requirements set by the human hearing system in order to avoid audible artifacts caused by digitizing the analog signal. For example, the minimal sampling rate needed is about 350 kHz, which is not covered by a 192 kHz PCM recording. Moreover, from 20 kHz and onwards the necessary signal-to-noise ratio becomes increasingly less important, which is in concordance with the natural behaviour of a SDM. Due to the latter feature of DSD, the signal becomes very bit-efficient compared to existing PCM formats closest to the afore mentioned minimal requirements. These statements include lossless coding for both kinds of signals. For example, on SACD a 6-channel DSD recording of 95 minutes can be stored, whereas for 20-bit, 192 kHz PCM only 55 minutes can be stored.

#### REFERENCES

- [1] Adams, R. W., Ferguson, P.F., Ganesan, A., Vincelette, S., Volpe, A., and Libert, R., 'Theory and practical implementation of a fifth-order sigma-delta A/D convertor', *J. Audio Eng. Soc.* **39**, 515-528 (1991).
- [2] Angus, J.A., and Casey, N.M., 'Filtering  $\Sigma\Delta$  signals directly', 102nd AES Convention, Mar 22-25, Munich, preprint 4445 (1997).
- [3] Angus, J.A., and Draper, S. 'An improved method for directly filtering  $\Sigma\Delta$  audio signals', 104th AES Convention, Amsterdam, may 16-19, 1998, preprint 4737.
- [4] Baird, R.T., and Fiez, T.S., 'Improved Delta Sigma DAC linearity using data weighted averaging', *Proc IEEE Int. Symp. Circuits Sys.* **1**, 13-16 (1995).

- [5] Bruekers, F., Oomen, W., and Vleuten, R. v.d., 'Improved Lossless Coding of 1-bit Audio Signals', 103<sup>rd</sup> AES Convention, Sept 26-29, New York (1997).
- [6] Carbone, P., and Petri, D. 'Effect of additive dither on the resolution of ideal quantizers', *IEEE Trans. Instrum. Meas.* **43**, 389-396 (1994).
- [7] Dunn, J., 'Anti-alias and anti-image filtering: The benefits of 96kHz sampling rate formats for those who cannot hear above 20 kHz.' 104<sup>th</sup> AES convention, may 16-19, 1998, preprint 4734. Download available from <http://www.nanophon.com> (2000).
- [8] Eastty, P.C., Sleight, C., and Thorpe, P.D., 'Research on cascadable filtering, equalisation, gain control, and mixing of 1-bit signals for professional audio applications', 102<sup>nd</sup> AES convention, march 22-25, Munich (1997), preprint 4444
- [9] Galton, I., 'Granular Quantization Noise in a Class of Delta-Sigma Modulators', *IEEE Trans. Inf. Theory* **40**, 848-859 (1994).
- [10] Gray, R.M., and Stockham, T.G., 'Dithered quantizers', *IEEE Trans. Inform. Theory* **39**, 805-812 (1993).
- [11] Jager, F. de, 'Delta modulation - a method of PCM transmission using the one unit code', *Philips Res. Rep.* **7**, 442-466 (1952).
- [12] Hawksford, M.O.J., 'High definition digital audio in 3-dimensional sound reproduction', 103<sup>rd</sup> AES convention, New York, preprint 4560 (1997), and Hawksford, M.O.J., 'Bitstream versus PCM debate for High-Density Compact Disc', private publication (1995 April) available for download at <http://www.meridian-audio.com/ara>.
- [13] Lipshitz, S.P., Wannamaker, R.A., and Vanderkooy, J., 'Quantization and dither: a theoretical survey', *J. Audio Eng. Soc.* **40**, 355-375 (1992).
- [14] Lipshitz, S.P., and Vanderkooy, J., 'Why Professional 1-Bit Sigma-Delta Conversion is a Bad Idea', 109<sup>th</sup> AES Convention, Los Angeles, sept. 22-25, preprint 5188 (2000).
- [15] Norsworthy, S.R., 'Effective dithering of sigma-delta modulators', *IEEE Proc. ISCAS'92*, **3**, 1304-1307 (1992).
- [16] Norsworthy, S.R., Schreier, R., and Temes, G.C., *Eds.* 'Delta-Sigma Data Convertors', IEEE Press, 1997.
- [17] Oohashi, T., Nishina, E., Honda, M., Yonekura, Y., Fuwamoto, Y., Kawai, N., Maekawa, T., Nakamura, S., Fukuyama, H., and Shibasaki, H., 'Inaudible high-frequency sounds affect brain activity: hypersonic effect', *J. Neurophysiology* **83**(6), 3548-3558 (2000).
- [18] Vanderkooy, J., and Lipshitz, S.P., 'Dither in digital audio', *J. Audio Eng. Soc.* **35**, 966-975 (1987).
- [19] For example, the CS4396 DA from Crystal.
- [20] Brochure on DSD editing using the Pyramix editing system. Available from Merging Technologies ([ccellier@merging.com](mailto:ccellier@merging.com)) and Philips ([Kim.van.Erp@philips.com](mailto:Kim.van.Erp@philips.com)).
- [21] The Sonoma editing system has been demonstrated at the 108<sup>th</sup> and 109<sup>th</sup> AES conventions.