Technical Details of the Sabre Audio DAC

Martin Mallinson and Dustin Forman, *ESS Technology Technical Staff*

Abstract—**The Sabre DAC is an octal audio digital to analog converter each individual channel of which performs with a DNR of >127dB. If used as a Stereo DAC the performance exceeds 130dB. The THD of each of the eight DACs is less than 0.0001% (120dB). This technical note outlines the critical components of the Sabre DAC that contribute to these performance levels.**

I. INTRODUCTION

THE Sabre DAC is an eight channel audio DAC designed to provide eight channels of high quality audio "surround sound" analog signals or two channels of exceptionally high quality Stereo analog signals. In the Stereo application the Sabre DAC provides a Dynamic Range (DNR) in excess of 130dB[1]. Multiple digital interfaces are provided: four I^2S interfaces operating in slave mode (also accepting all the various justifications) provide data to all eight channels; DSD[2] is also accepted for all eight channels and finally a built-in SPDIF[3] decoder is provided for the Stereo mode. The Sabre DAC operates from nominally 3.3v for its analog section and nominally 1.2v for its digital section. The analog section power supply is flexible: it will operate from 4.0v down to1.8v with little performance degradation. At the nominal 3.3v analog supply the Sabre DAC consumes a total of 120mW for data rates up 48KS/s (160mW at 192KS/s) and at the lower 1.8v analog supply consumes 65mW for data rates up to 48KS/s (100mW at 192KS/s)[4].

ESS Technology has shipped millions of audio products operating in the >100dB level as part of complete DVD chips and it is lessons learned in the harsh environment of multi-million gate chips that we are able to leverage to create the high performance of our stand-alone Sabre DAC products.

II. OVERVIEW

The details are given in the following sections, but in overview, the performance of the Sabre DAC is achieved by minimizing and simplifying the analog components on the chip – there is not a single analog bias line in the entire design, there is no closed loop analog element and no MOSFET device operates in other than the "off" and "saturated on" region. There is no analog PLL: experience shows that even a carefully designed PLL is prone to interference from the digital section and modulation of its phase noise in any way related to the signal is detectable by the human ear. In addition to the simplified analog section, care is taken in the digital signal processing (DSP) to ensure that jitter rejection is essentially perfect: a patented technique is used to re-create the audio data in a crystal-controlled low phase-noise clock domain completely isolated from the clock domain of the transport medium and so not at all related to the clock domain in which the data was sampled. This latter feature ensures that there is zero jitter transfer from the various interface clocks to the digital data and that the digital clock of the DAC itself can be optimized for noise and power consumption[5]. Finally, operation of the DSP in the transport medium clock domain (the "Fs" domain in which data is delivered to our DAC) is done in a "gated clock" mode – a form of serial/parallel processing that prevents any inter-modulation of operating current between the DAC clock and the data transport clock that could cause harmonically related noise on the system power supplies. The following sections explain these features in more detail.

III. EXTRACTING THE DATA FROM THE TRANSPORT MEDIUM

A. SPDIF Interface

The SPDIF interface is more complex than the DSD and I^2S since it must first derive the embedded clock in the bi-phase encoded data. In fact, experience with many forms of SPDIF decoder suggest that most fail in the presence of high jitter due to the lack of robustness in the clock recovery process. To avoid this potential problem the Sabre SPDIF interface avoids having to extract the clock at all: decoding is done using a method that does not require an explicit measure of the clock frequency. Specifically, the digital input[6] is first corrected for

[1] DNR as measured by the Audio Precision 2722 using the 22Hz High Pass, 22k Low Pass and A-weighted filters.

[2] DSD is Direct Stream Digital and is the data format used by SACD – Super Audio Compact Discs.

[3] SPDIF (or S/PDIF) is the Sony-Philips Digital Interconnect Format which is a Bi-Phase encoded data stream of stereo audio samples conveniently using only one signal wire and common to many audio component systems.

[4] These power consumption figures are significantly lower than competitive devices and include all eight channels running concurrently.

[5] The DSP engine of the Sabre™ DAC operates at either 27MHz for audio rates up to 96KS/s or 40Mhz for audio data rates up to 192KS/s. Neither of these rates is integer-related to the audio clock.

[6] The SPDIF input is a conventional CMOS input pin – it is not a comparator – the user must provide a comparator prior to the input pin if it is needed.

50% duty cycle by means of a discrete digital delay line that is able to delay either the positive edge or the negative edge of the signal such that after this delay line the signal is at 50% duty cycle. Thereafter an assessment is made of the width of each pulse based on its relation to recently seen pulse widths and a decision circuit assigns each a width of 1, 2, or 3 units. A state machine then operates on the assigned widths in succession; this state machine is searching for the block boundaries and the bit states. The state machine makes no attempt to re-time or otherwise decode the clock – it simply "time stamps" the event and passes it to the downstream processor. Using this method the SPDIF interface is able to accommodate 50nS of random jitter and 200nS[7] of sinusoidal jitter in the incoming data.

B. Rendering the Data into the DAC clock domain

The I²S and DSD are clocked data streams, they do not require the method used by the SPDIF interface, but they result in the same thing: similar to the output of the SPDIF decoder, at a certain asynchronous time, they present to the downstream processor the new bit or word from the interface. Rather than trying to lock a PLL to the data rate of this interface signal, the Sabre DSP uses the arrival time of the data as a gating signal for the first part of the processing. Specifically, we now need to apply a filter to the digital data – we must remove the image that will be created when we up-sample the data to our much higher clock rate. This filter involves a number of cycles of the DSP. One desirable consequence of this method is that the over-sampling filter time constant tracks the data rate. The data we have after this process is, of course, mathematically correct, but if we use this data in the higher clock domain we will have a great deal of noise, since as soon as we try to sample the data we must decide at which edge of our high speed clock this data is to be used. That choice is never correct – this sample delivered from the interface is supposed to be at some point between our high speed clock edges. The problem then is this: we have the data stripped off the transport medium – the data is mathematically correct but we don't know where in time this data is supposed to be, and even if we did we cannot snap it to our high speed clock because that represents a quantization in time and hence noise. The conventional solution is well known: first a digital PLL is used to remove the jitter of the incoming data (since it will suffer from jitter in the transport clock) and then a poly-phase filter is employed to rate-convert the signal[8] into the new clock domain. These kinds of sample rate converters work well but they are limited to certain ranges of operation (they have a limited ratio of rate conversion – typically about 8:1) and a DNR less than that in the digital data itself. The Sabre DAC rate converter has two advantages compared to the poly-phase filter approach and is described in great detail in the pending patent. Firstly the rate conversion is unlimited – allowing the Sabre to always achieve a conversion into its

exceptionally high clock rate (as much as 40Mhz) from as little as 4Khz in one step; and secondly, the process is essentially perfect to the bit level – the output DNR exceeds 175dB and the THD is correspondingly high.

As well as sample rate conversion the Sabre has a proprietary jitter reduction circuit that operates with the rate converter and is able to achieve a 100% jitter rejection. These two steps: jitter rejection and rate conversion; are able to take the "burst" mode over-sampled filter output into the precisely correct clock edge of the high speed clock. Audio data from all sources is now in the high speed clock domain and sent to the modulator.

IV. THE HYPERSTREAM™ MODULATOR

ESS uses the trademark HyperStream to refer to our range of noise shaping, bit-reducing modulators[9]. Our modulators differ in detail from conventional $\Sigma\Delta$ modulators in that they can operate up to 100% modulation depth. This is achieved by a cascade of independently stabilized lower-order modulators and careful selection of relative gain in the integrator sections such that the onset of clipping is well-controlled. As the modulation depth approaches 100% the loop order essentially softly degrades to a lower and lower order, each lower order is independently stable hence the whole modulator is stable for any modulation depth. This allows us to operate with a nominal full scale of 90% modulation depth in the Sabre DAC multi-bit modulator[10]. The modulator is fifth order with adjustable (user software configurable) output bit width, typically six bit, and since it operates to 90% modulation depth and at rates up to 40Mhz, the noise floor from the modulator is well below −160dB (the integrated noise is less than −140dB).

Because $\Sigma\Delta$ modulators are very non-linear systems only approximations of performance are possible in frequency domain simulations. Time domain simulations are useful to verify first-order stability (i.e. that the loop does not oscillate in the typical condition) but beyond this, simulations must be augmented with hardware tests. We have performed extensive hardware testing of the HyperStream modulators over many years (and many millions of shipped products) and consequently have designed them to show additional features that are not commonly tested. We believe that some of these features of the modulator are the reasons why our users tell us that the Sabre "sounds better" than similarly specified parts. Two of these features are explained in detail here.

Firstly, it is a known problem in noise shaping modulator design that noise is not independent of DC signal level. A graph of DNR vs. DC signal level will always rise as the modulation depth approaches 100%. Most modulators cannot handle 100% modulation depth and what they call full-scale may be only 50% modulation depth, nevertheless a graph of DNR vs. modulation depth up to even this 50% level will show an increasing noise. The HyperStream modulators are

---

[7] 200nS is even greater than the nominal SPDIF clock period – but nevertheless the SPDIF interface correctly decodes the data stream if its FM modulation is at a sufficiently low frequency.

[8] For an introduction to this technique see the Analog Devices AD1896.

[9] To understand how ESS modulators may differ from more conventional $\Sigma\Delta$ modulators refer to US Patent 7,058,464 which describes one difference. Other innovations are not publicly disclosed – they are retained as ESS Proprietary information.

[10] All other things being equal, this 90% modulation depth alone delivers a 5.1dB increase in DNR.

designed such that this increasing noise is not so pronounced even as 100% modulation depth is approached (the modulator on the Sabre DAC has full-scale set to 90% modulation depth and even so shows less degradation of DNR on this test). But worse than this, a careful inspection of noise vs. DC offset will show that the noise is non-monotonic: some DACs show noise spikes that degrade the DNR to as much as 30dB less than the specification at certain specific DC offsets. The Sabre DAC shows no more than −123dB DNR even at the worst case DC offset.[11] A problem similar to noise vs. DC offset can show itself when noise vs. rate-of-change of DC offset is plotted. In fact any constant derivative of input may cause an anomalous noise; the Sabre is designed to minimize this phenomena as well as the simple DC offset case.

Secondly, certain $\Sigma\Delta$ modulators when provided with a rapidly changing input signal will exhibit non-linear noise behavior as they process the transient. This is because all noise shaping modulators are feedback systems and the usual design process (supported by commonly available design tools) operates to minimize in-band noise suppression while maintaining stability. This noise-optimized stable loop configuration will lead to an output that matches the input to the required degree within the requested bandwidth as expected. However this typical design process neglects the dynamic response of each state variable: there are choices of Q (and relative gain) that minimize noise but result in relatively large lightly damped resonances of the internal state variables. The consequence of this is that in a quiet passage of music the state variables of the modulator are all operating within a certain "state space" and the quantization noise shaping is described by the noise characteristics in this "volume" of the space. After a large music transient has passed, the output traces its dynamic response back to the quiescent operating point as we expect, but every state variable is also following its transient response back to its quiescent point[12]. During this multi-dimensional excursion back to the lower signal level the operating point traverses different volumes of the space, each of which has its own noise characteristic. Hence a very perceptive listener can hear something "anomalous" related to the transient response. We have designed our HyperStream modulators to exhibit strongly damped responses in all state variables. This means that a low level signal processed just prior to a transient and just after a transient, is processed in the same corner of the state space of the modulator and hence in the same quantization noise environment, thereby eliminating the anomalous errors in the music reproduction. These are the intrinsic reasons that the Sabre DAC reproduces music so accurately.

There are other benefits from our work on HyperStream modulators within our DVD chips which are not mentioned in detail here but they relate to the need to have all digital noise sources (truncation, clipping) give rise to quantization error signals that have precisely zero average value. The Sabre modulator has all these features.

V. THE ANALOG DAC

After all the DSP and complex noise shaping of the signal is complete the digital number must be converted to an analog output. In principle the typically six bit number may be applied to a six bit DAC and the analog output is produced. However, without additional sophistication, that analog output will never be more accurate than the DAC that is used – so the design is much more complex than a simple connection to a six bit DAC. This DAC was the focus of a great deal of effort in the chip development: the connection of the elements that make up the DAC to the analog power supply is exquisitely sensitive to any mismatch[13]. However, to achieve a reference that is sufficiently low noise, to not compromise overall performance, requires an external voltage reference (AVcc). Therefore, the node labeled AVcc (the analog power supply) might better be called "DACReference" since it is the DAC reference level. You will note in our recommended PCB designs that either an amplifier or a large capacitor is placed at the Avcc point to ensure that the noise is not reduced by any spurious signal at this point. Further insight into the critical nature of the DAC design may be gained by noting that the crosstalk (from one DAC to another in Stereo mode) is less than −135dB – a testament to the degree of care taken in the on-chip routing of signals related to the DACs.

Despite all this effort in chip layout the DAC still would not meet the specifications were it not for a further innovation: a patented form of Dynamic Element Matching[14] (DEM) that is able to frequency shift any residual DAC mismatch well out of the audio band. Details may be found in the patent, but basically this particular form of DEM maintains a high degree of error cancellation such that in-band tones or spurious signals are completely removed.

One further aspect is worth mentioning: no matter how accurately the DAC is constructed, jitter will degrade the fidelity of the audio signal: if any jitter, that is time domain uncertainty, is associated with the data path that brings the bit from the digital DSP engine to the DAC, that jitter will show itself as a significant additional noise[15]. The DEM and specific switching pattern applied at the DAC bit is optimized to reduce sensitivity to timing jitter, but to mitigate this effect still further the distribution of the clock from the lowest phase noise point (which is the output of the on-chip crystal oscillator itself) to the sampling gate of the D-Type that times the signal into the CMOS inverter of the DAC, is done with extreme care. Each gate is optimized for maximum power supply rejection; each clock line is carefully laid out to remove any data dependant modulation of timing interval and the intrinsic timing match on the dedicated clock path is held to the sub-picosecond level.

The Sabre DAC has a total of eight such DAC channels, but each DAC is actually a balanced pair of the DACs as

---

[13] The design achieves a systematic error of less than $700\mu\Omega$

[14] US Patent 7,116,257

[15] The noise that jitter induces is not easily described: it is not a harmonic distortion but is a noise near the tone of the music that varies with the music: it is a noise that surrounds each frequency present in the audio signal and is proportional to it. Jitter noise is therefore subtle and will not be heard in the silence between audio programs. Experienced listeners will perceive it as a lack of clarity in the sound field or as a faint noise that accompanies the otherwise well defined quieter elements of the audio program.

---

[11] The Audio Precision 2722 can perform this test. Use a slowly changing DC offset superimposed on the -60db 1k tone to see the noise artifacts.

[12] The operating point in the machine state space is returning to the quiescent position – all the dimensions of the state are changing.

described, so there are in effect 16 such DACs. Considerable flexibility in bit width and channel assignment of the DACs is provided, but the default configuration is optimum and unless some special need is found we recommend use of the default configuration.

The default configuration, in eight channel mode, is that a pair of six bit DACs operating in anti-phase make up each channel. Since the DACs are as described, namely a Thevenin equivalent to a voltage source linearly between AGnd and AVcc having an approximately 800Ω output impedance, you may choose to connect this output to a voltage mode (non-inverting amplifier) configuration; or you may choose to connect it to a virtual ground current mode (inverting amplifier) configuration. The highest performance in terms of THD is via the current mode[16], but both voltage and current mode provide about the same DNR.

The default configuration for Stereo mode, and the only configuration that can use the SPDIF input, is to wire four output channels in parallel. When Stereo mode is enabled in the configuration registers the same data is sent to all four channels. Effectively now the DACs become a pair of eight bit DACs having an output impedance of about 200Ω. This configuration allows >132dB of DNR to be typically reached[17]. THD in the Stereo current mode is limited by the external components and measurement equipment. We recommend using an extremely good op-amp for the highest performance but even an excellent op-amp is the limiting factor in the THD[18].

VI. Other Considerations

Aside from the items so far discussed, the Sabre DAC has these additional aspects:

A. Very Wide Data Path

The Sabre DAC arithmetic operations are done with a wide data path width: 48 bits of precision are maintained in the critical element of the DSP. This extraordinary high resolution is then truncated with a noise shaping of its own such that the resulting digital number fed to the HyperStream modulator has less than −170dB of digital noise due to the calculations performed in the DSP. This emphasizes the fact that, to truly appreciate the Sabre DAC performance, you must use at least a 24 bit audio data source.

B. Zero Stuffing not First Order Hold in up-sample

A first order hold in the over sampling filter results in a narrower data path and hence less logic, but necessitates sin(x)/x correction be at a lower rate. No sin(x)/x correction is perfect and low rate sin(x)/x correction results in frequency domain imperfections. Zero-stuffing is mathematically more precise and results in a sin(x)/x correction at a much higher frequency and hence any errors are much further out-of-band.

C. No Peaking in the Linear Phase Filter

The linear phase filter used to up-sample the digital data is not used to compensate any lack of flatness in the succeeding sections. Although an apparently simple means to correct for lack of gain flatness, this technique leads to slight peaking in the digital domain that disrupts the phase linearity of the output and can show up as unexpected variation in output group delay vs. frequency[19]. Consequently, the Sabre DAC has independently flat modulator and filter configurations.

D. Clickless, Very High Resolution Volume control

The firmware presents a 0.5dB volume control register to the user allowing setting of volume to 0.5dB accuracy. However, internally the volume control operates to less than 1/64 of a dB and smoothly moves from one requested volume to the next. This results in a perfectly click-less volume step.

VII. Conclusion

The Sabre DAC is a collection of circuits, techniques and systems designed to make the best DAC ever built. One feature that allows the Sabre DAC to reach such high performance is the modulator design: listening tests show that experienced listeners can distinguish between a conventional ΣΔ modulator and the extensions to the ΣΔ modulator as implemented in our range of HyperStream modulators. Experiments performed in our listening room suggest that it is the absence of correlations in parameters that seem to contribute to the listening experience. Specifically, two potential correlations are important: the signal level and noise level do not correlate – noise does not perceptibly vary with amplitude or rate of change of the signal; and transient events are rapidly extinguished in the integrators of the modulator which suppresses anomalous noise potentially created as the transient passes through the signal path.

However, to expose this level of detail to the listener, great care must taken in preventing jitter and suppressing any mechanism of cross talk between the data and the phase noise of the system clock.

Example Performance Graphs

---

[16] Current mode cancels a slight on-chip resistor voltage coefficient.

[17] With 4.0V power supply greater than 133dB can be achieved.

[18] This is demonstrated by reducing the amplifier output signal excursions: when this is done DNR is slightly reduced but the THD improves, thus demonstrating that it is the amplifier that is limiting the THD. Beyond about 124dB THD limitation of the Audio Precision 2722 that is used to characterize the Sabre DAC begins to become apparent.

[19] Specifically, the digital filter is a symmetric FIR filter and has linear phase. Any residual frequency domain anomaly in the modulator is probably not linear phase; hence correcting the frequency domain response by adjusting the FIR appears correct in the frequency response plot, but if group delay is plotted it will be seen to be variable and possibly non-monotonic. (Essentially this is stating the obvious fact that group delay variation cannot be corrected in a FIR that has linear phase – the modulator and filter must be designed independently, neither relying on the other for any correction).

Please refer to the ES9006, ES9007 and ES9008 data sheets for details of the performance expectations from each member of the family of Sabre DAC parts. This technical document has focused on the ES9008.

Some FFT data from the ES9008 (captured on the AP 2722 system) are shown on the right. The upper plot (Figure 1) shows the THD. Note that the highest harmonic is the 5th and is at about −120dB. (As mentioned in this note, this harmonic is due to the external amplifiers). The middle plot (Figure 2) shows the DNR performance with −60dB input signal.

The lower plot (Figure 3) compares the jitter performance of a typical competitor's part (red, or upper line) with the Sabre DAC (blue, or lower line). The graph shows THD+Noise vs. Frequency with only 2nS of random jitter on the data transport clock (on the clock of the I^{2}C in this example).
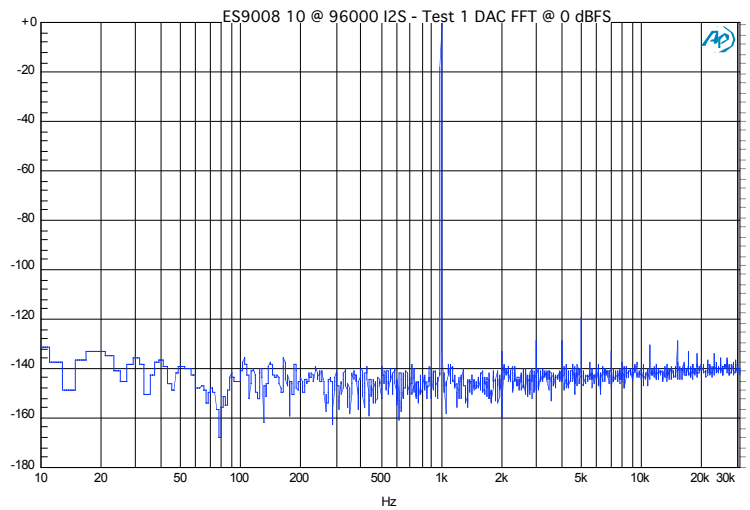
ACKNOWLEDGMENT

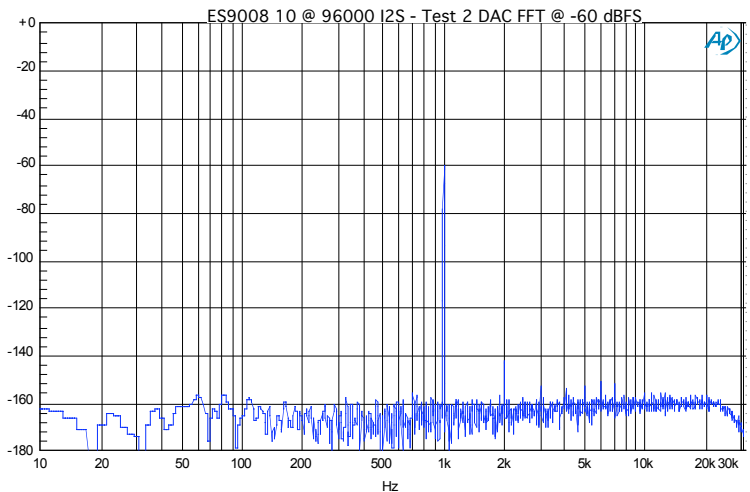Figure 1: FFT with 0db input signal showing THD
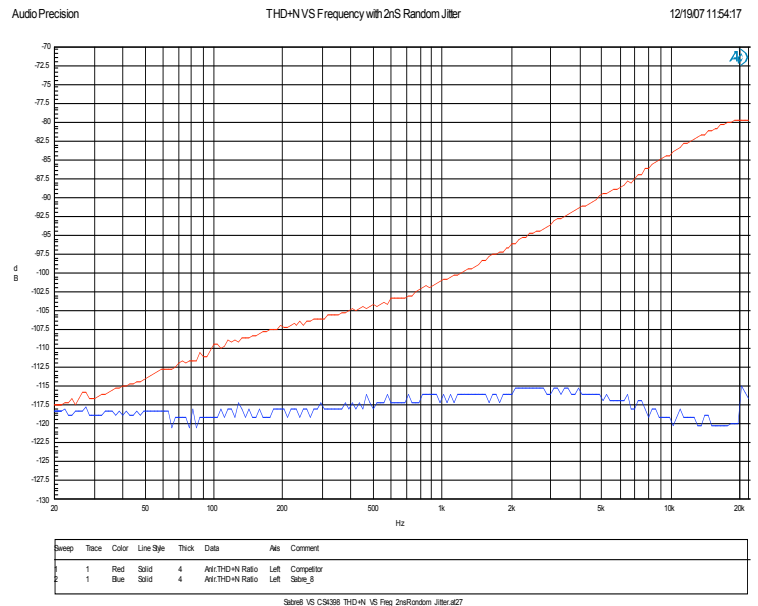

Figure 2: FFT with input at -60db showing DNR

Figure 3: The effect of 2nS Jitter on THD+Noise vs. Frequency. The upper line is a typical competitor's part and the lower line is the Sabre DAC.