Audio Engineering Society

Convention Paper

Presented at the 112th Convention
2002 May 10–13 Munich, Germany

A True One-Bit Power D/A Converter

Bruno Putzeys[1] and Renaud de Saint Moulin[1]

[1] *Philips Digital Systems Labs, Leuven, B-3001, Belgium*
Correspondence should be addressed to corresponding author (bruno.putzeys@philips.com)

ABSTRACT

A method is presented to convert 1-bit digital audio signals into an analogue signal with sufficient current and voltage to drive loudspeakers. For this goal a novel non-PWM class D power stage is constructed that performs this function with very low distortion and very high efficiency, without the use of feedback or other analogue processing. Results of the prototype development are detailed.

INTRODUCTION

A continuous theme in audio amplification the past decade has been the desire to merge class D power stages with digital modulators, forming a "digital amplifier" or a "power DAC". While the authors believe this desire to be subjective in origin, the question is accepted "as is" without further elaboration. The most obvious setup is to generate a PWM (-alike) signal through DSP or other digital means. Another approach often looked into is directly to use 1-bit deltasigma signals such as what can be found in many digital-to-analogue conversion systems. Although the latter is fraught with difficulties normally considered insoluble, it has resurged due to the introduction of a 1-bit deltasigma based release format called DSD. The present method is capable of surmounting the audio quality and efficiency related problems of 1-bit power conversion to such an extent that it consistently outperforms PWM based methods.

PROBLEMS FOUND WITH 1-BIT SIGNALS AND REAL POWER STAGES

For practicality we shall use the term "1-bit" whenever a 1-bit noise-shaped signal at typically 2.8224MHz sampling rate is meant, such as DSD.

Switching Frequency

Spectral and mathematical analysis of a 1-bit signal shows that a strong discrete tone is present which is at idle exactly half the sampling rate (corresponding to frequent and long …101010… patterns) and which decreases linearly with modulation index[1]. This means that the effective repetition rate is around 1.4MHz, improving down to below 1MHz for strong modulation. In class D power circuits such high frequencies are to be avoided in order to keep switching losses acceptable and to reduce the contribution of timing errors to output distortion.

Switching Losses

In one-quadrant power systems such as DC/DC converters, switching losses can generally be minimized by the use of zero-voltage and/or zero-current switching techniques. These recover energy stored in parasitic circuit elements. In four-quadrant mode these energies can no longer always be recovered and must be dissipated inside the active elements (MOSFETs). Additionally, the internal diodes of the power MOSFETs will sometimes be forward biased, incurring very significant losses during forced turnoff. These issues dictate a need for minimizing switching rates.

Distortion

When the output stage has settled, it has firm control over the output voltage. While transitioning this is not the case and the actual voltage waveform found will depend upon the output current delivered. Attempts to correct for these errors in a digital control circuit without feedback must inexorably fail as the load impedance and hence the output current is unknown.

The error committed under large signal conditions by an open-loop class D power stage due to a timing uncertainty t_d (roughly equal to dead time) is approximately a square wave of amplitude $t_d \cdot f_{sw} \cdot V_{cc}$ i.e. the supply voltage times the fraction of time spent in transition. The harmonic content of a square wave is 0.41 so the error can be estimated as

$$0.4 \cdot t_d \cdot f_{sw}$$

If the power stage has a long dead time to employ the natural ZVS regime that occurs at low signal levels (only with binary modulation), THD at medium power levels degrades linearly with signal level i.e. the absolute THD level remains roughly constant outside the ZVS regime.

With shorter dead times (and hence shorter rise times) the medium-level situation may improve by as much as 10dB ("rounding off" of the square wave's corners). Unless very low idle losses are required, the latter option is usually chosen. Still, the proportionality with timing uncertainty and switching frequency remains.

A power stage driven by a 1-bit signal with its high repetition rate spends more time transitioning and thus commits greater errors. These errors not only show up as distortion, but also as demodulated shaped noise.

Modulation Index

As a 1-bit signal is modulated further, its repetition rate decreases. The relative rarity of either 1's or 0's in the data stream (depending upon the polarity of the excursion) implies lower information content and hence less bandwidth for the noise shaper to shape noise into. At some point the noise shaper will become unstable. The maximum stable modulation level is a factor when designing a noise shaper and it is typically set at 3 to 6 dB below "rail-to-rail". The DSD reference "0dB" level is defined at 50% modulation index.

This limits the maximum available output power to one quarter of what the power stage is physically capable of. Conversely, for a given output power, supply voltage must be up-rated to 200%, making for bigger and slower FETs with larger parasitics and much more recovery charge. This further exacerbates switching losses.

Summed Up

With a traditional power stage, 1-bit signals are the least efficacious way of driving it. A typical open-loop design would have no better than 70% efficiency, 1% THD and 70dB SNR. One company is producing amplifiers in which the audio performance is corrected by placing the power stage inside an analogue control loop, with an analogue input. This goes against the premise of a "digital amplifier" and the efficiency issue remains unsolved.

MULTIPHASE POWER STAGES

Multiphase power stages are becoming common practice in the DC-DC converter business (more precisely for CPU supplies) and have also been described for audio[2]. Here, *n* small power stages (each of a silicon area 1/*n* compared to a single power stage) are combined using *n* output inductors into one.

The main advantages of such a construction lies in the reduction of the physical switching frequency (per phase) by a factor of *n* with full retention of the bandwidth. This will reduce switching losses to insignificance and reduces distortion in two ways, namely by reducing the switching time (as allowed by the smaller MOSFETs) and by reducing the significance of the error by a factor *n*.

Shown is the original circuit proposed under [2].

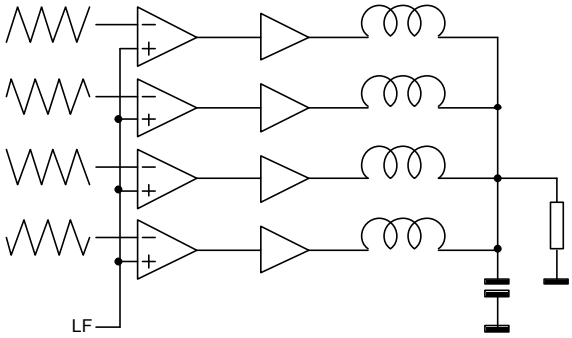

This particular topology is difficult to use, for two main reasons.

- While the modulation byproducts cancel out at the summation, they can be quite significant when the signal frequency is high and the carrier frequency is low, and will appear as circulation currents. Other timing errors will also create circulating currents. A minute DC error between legs can easily lead to saturation in the corresponding pair of inductors.
- To obtain good summation precision, required for open-loop operation, the coils need to be matched to DAC-like precision.

These two problems will be worked one by one.

THE SAMPLED MULTIPHASE POWER STAGE

An alternative way of driving an *n*-phase output stage is to view it as an *n*+1 quantum D/A converter, and feeding it *n*+1 level data words.

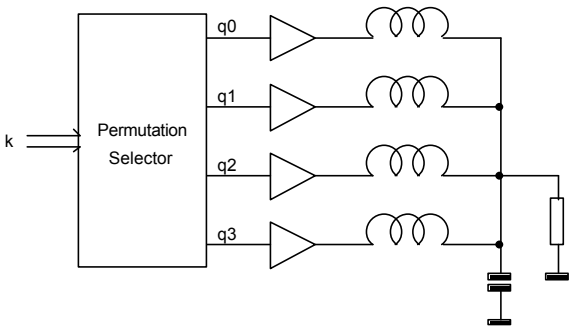

To control the unwanted circulating current in a sampled-time system, a measure of the circulating currents is obtained by digitally integrating the "voltage" (in reality the number of quanta) found across each inductor over sampled time. For a given data value *k*, the number of ways to arrive at this value is

$$\frac{n!}{k! \cdot (n-k)!}$$

allowing a "permutation selector" to choose, at any given instant, the permutation which results in a desirable evolution in circulating currents based on the estimated currents found, in compromise with a low switching frequency on each leg. The latter is achieved by requiring that the Hamming distance between two consecutive drive words *q* be equal to the difference between the two corresponding *k* values except when one of the coils is coming close to saturation in which case an extra pair of compensating bit flips is allowed.

Another way of seeing it is that in order to keep the inductor current within bounds, the voltage across it must have zero DC component. Spectrally the difference signal found between any two switching legs will be an "empty" first-order noise shaped signal.

This method effectively controls the circulating currents while allowing drive by an *n*-level noise shaped signal. Such a signal can either be made directly or derived from a 1-bit signal.

Trivial conversion of 1-bit data into a suitable drive signal

The most intuitive way of viewing the control of the multiphase power stage by a 1-bit signal is as *n* consecutive 1-bit samples out of a shift register being assigned via a permutation selector to the *n* phases. The state of the output stage corresponds at all times to the number of ones in the current time window. No decimation occurs: the shift register and the output stage are updated every time a new sample comes in. Desirable characteristics of this method are:

- No addition of spectral components (such as through reshaping)
- General low-pass characteristic caused by the rectangular impulse response, lessening the requirements on the output filter.
- The length of the impulse response is even. This produces a spectral zero at half the sample rate, dramatically reducing the number of transitions in the converted signal compared to the original 1-bit signal.

Extension Of The Modulation Index To 100%

It was said earlier that a 1-bit signal will typically be modulated only to 50%. After conversion to an n+1 level signal this means that only the (n/2)+1 middle levels will effectively be used. This is compensated for by extending the "bit history" by a factor two, resulting in a 2n+1 level signal of which only the n+1 middle ones are used, making full use of the power stage modulation capability.

It will be obvious that due to the out-of-band noise the modulation will momentarily exceed the n+1 levels from time to time. Experimentation has shown that by "holding" excess quanta and filling them in later when the modulation allows, full 100% modulation can be achieved at only a minor THD penalty over the top 1dB of modulation for an 8-phase power stage. This is a first-order noise-shaping action, really, which only takes place in the event of clipping.

INDUCTIVE SUMMATION OF PHASES

Until now the summation of the phases was considered as implemented using multiple output chokes. It is understood that a multilevel DAC driven by a noise-shaped signal will need to have its output levels fixed to a precision commensurate with the THD and the s/n intended. With a 100dB spec this requires a 0.001% tolerance from the side of the inductors. We will not be making an amplifier in this manner.

The solution lies in splitting the combination and filtering function into separate blocks. The output filter coil will again be a normal coil and a device to combine two phases will look like this:

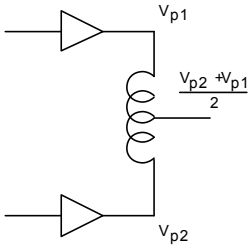

On the prototype, small toroids were used with the two winding halves bifilar wound to insure optimal coupling. The load current is shared between the two phases and will see only the leakage inductance of the coil (i.e. the magnetic fields caused by the two halves of the output current cancel out). A practical value of about 20nH is found. The inductance as seen between the two ends of the coil is large (typically 1.8mH on the prototype), supporting only a small circulation current which is kept under control by a modified permutation selector. An upper bound to the precision of the combination can be found as follows: if there is an error x between the inductances of both winding halves, it should show up as leakage inductance. This implies that x must be less than or equal to the ratio of the leakage inductance to the total inductance, which is about 0.0011%.

Tree Summation

For this development it was decided to build an 8-phase power stage where the load would be tied in bridge between two 4-phase "half bridges". Combination of four phases looks like this:

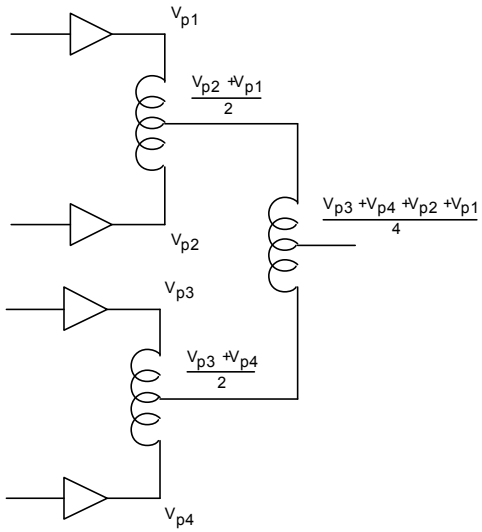

It can be seen that in this design, the multilevel signal will be physically present before the output filter.

A PRACTICAL DESIGN

The block diagram of a complete design is shown here:

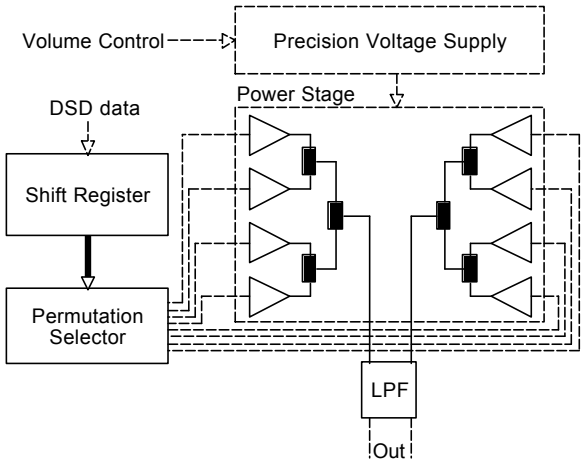

Each phase is a dual MOSFET in an SO-8 package together with a half-bridge driver, also in an SO-8. The switching rate is about 100kHz so the contribution of switching losses to heat dissipation may be neglected, counting only resistive losses. The on-resistance of the FETs is on the order of 150 mΩ (when hot). At a conservative rating of 1W per package the combined output current may be $10A_{rms}$. The power rating is 150W into 8Ω or $34.6V_{rms}$

so the power stage is effectively capable of holding its own up to 350W continuously into 3.4Ω, at which point the power devices are collectively dissipating 8 watts. This translates into a power stage efficiency of 97.7%. Thermal checks on the prototypes confirm this estimate.

The switching speed of the phases is less than 2ns with the timing error in the order of 500ps. According to the formula given earlier, distortion would be estimated to be on the order of 500ps·100kHz·0.4=0.002%. While this may be true for the power stage alone, total performance will also depend upon the output impedance of the power supply. The THD caused by the power supply's output impedance in a full-bridge open-loop class D power stage is approximately given by:

$$\mathrm{THD} = \frac{Z_{\mathrm{psu}}}{4 \cdot Z_{\mathrm{load}}}$$

The output impedance of the power supply regulator is limited by the ESR of a capacitor bank (8 times 1000µF/63V, each installed on top of one phase) which totals 3mΩ. Below the frequency given by the RC time constant of the elcaps (6.6kHz), the power supply regulator will provide a lower impedance. At frequencies above 16kHz, the parasitic inductance of the capacitors will dominate.

At 3mΩ, the combined distortion into an 8Ω load should be on the order of 0.01%.

PERFORMANCE MEASUREMENTS

The prototype was wired to an Audio Precision test set with a PCM-to-DSD converter to allow connection to the AP's digital outs. The power supply was set to 52V and an 8Ω load attached.

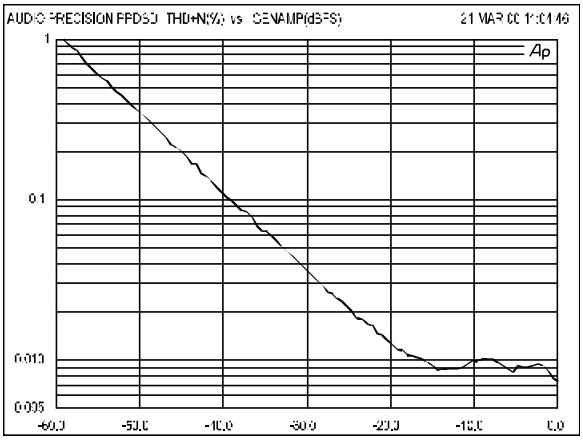

The THD versus power plot is dominated by noise up to about -20dB where distortion products up to about 0.01% become visible. Unweighted signal-to-noise ratio is 99dB.

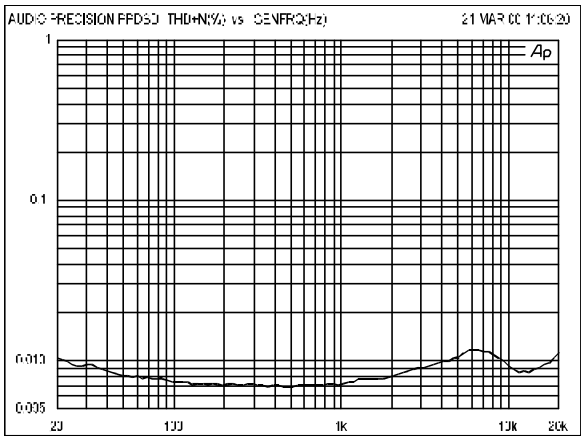

The THD versus frequency plot shows 3 mechanisms in action. At low frequencies, thermal modulation of the MOSFETs' on-resistance becomes obvious. Above 1kHz both the rise of the power supply impedance and some distortion from the output inductor show up.

SOUND QUALITY ASSESSMENT

The sound quality of the power D/A converter was tested both using normal CD material converted to 1-bit 64fs (DSD), and SACD material with the DSD signal directly fed into the converter. The sonic performance was found to be well in line with the expectations, i.e. on a par with the most exclusive audiophile analogue combinations.

CONCLUSIONS

The feasibility of direct power conversion of 1-bit deltasigma audio signals has been demonstrated. The initial drawbacks associated with 1-bit power conversion have been successfully overcome and actually turned into an advantage for the new topology when compared to traditional class D power stages.

REFERENCES

[1] John Vanderkooy & Stanley Lipshitz, "Towards a Better Understanding of 1-Bit Sigma-Delta Modulators," presented at the 110th Convention of the AES

[2] Karsten Nielsen, "Parallel Phase Shifted Carrier Pulse Width Modulation (PSCPWM) - A novel approach to switching power amplifier design," presented at the 102nd Convention of the AES.