PERFORMANCE CHALLENGES IN AUDIO CONVERTER DESIGN

IAN DENNIS

Prism Sound group, Cambridge, UK

The habitat of audio converters is ever-changing. Early converters were built into stand-alone digital audio equipment. Later, they were built as stand-alone conversion devices with dedicated digital audio connectivity. Nowadays they are increasingly offered as peripherals for PCs or computer networks. All the while performance targets have been rising. In some ways these changes have impacted audio converter design; in other ways it's business as usual. Herein are the random musings of an old warhorse, struggling to focus on a few key strategies in the eternal battle for optimum audio converter performance.

INTRODUCTION

It seems to me that the plan of campaign for most engineers setting out to design an audio converter subsystem (by which I mean a stand-alone converter or that part of an equipment which deals with audio conversion) is as follows: First, choose a data converter device (by which I mean an ADC, DAC or CODEC chip) which will meet the project requirements - performance, channel count, cost, features etc. Then implement a design around it based on the manufacturer's application note plus whatever additional features, interfaces etc. are required.

Yet the majority of converter subsystems perform somewhat below the potential of their chosen data converter – either all the time, or in certain situations. The random topics which follow are the result of my personal prejudices about why that usually is.

If you're wondering whether to read on, I can summarise the battle plan under these headings:

- Sort out your clocking;
- Learn to tame the switch-mode power supply;
- Strive for a top quality analogue signal path;
- Don't overlook the voltage reference;
- Make sure the digital parts don't ambush you.

But first, I should explain…

WHY DATA CONVERTERS ARE SORTED

When digital audio was new, the data converter itself was almost the only consideration – since it was impossible to build one with as much dynamic range and linearity as the professional or high-end consumer user was used to in analogue equipment. All we used to ask was 'what chip is in it?'

Nowadays, where workmanlike data converters can exceed a dynamic range of 130dB and THD+n of -110dB (audio band, rms, unweighted), the weak link in a conversion system is most often elsewhere. Surely there is no case for applying design effort/budget to the data converter itself, except in the most exacting of applications, and where it has already been applied in great measure to the rest of the converter subsystem?

OK, maybe the decimation and interpolation filters aren't sorted. Maybe there isn't enough silicon on the planet, I don't know, so I'm moving on. From here on in it's 'How to get the best out of your chosen data converter'. The mantra is 'Painstaking design and relentless assessment'.

DESIGN

Typical ADC and DAC subsystems

For clarification, typical ADC and DAC subsystems are shown in Figures 1 and 2. I've marked the parts which you have to take special care of (the analogue bits) and also the parts which will try to make that difficult (the rest). For spendthrifts, I've indicated a cut line where you might consider isolating the analogue and digital parts, but there are lots of ways you could do it, or not at all.

So, with the problem defined, here, in no particular order, here are my 'converter design' musings…

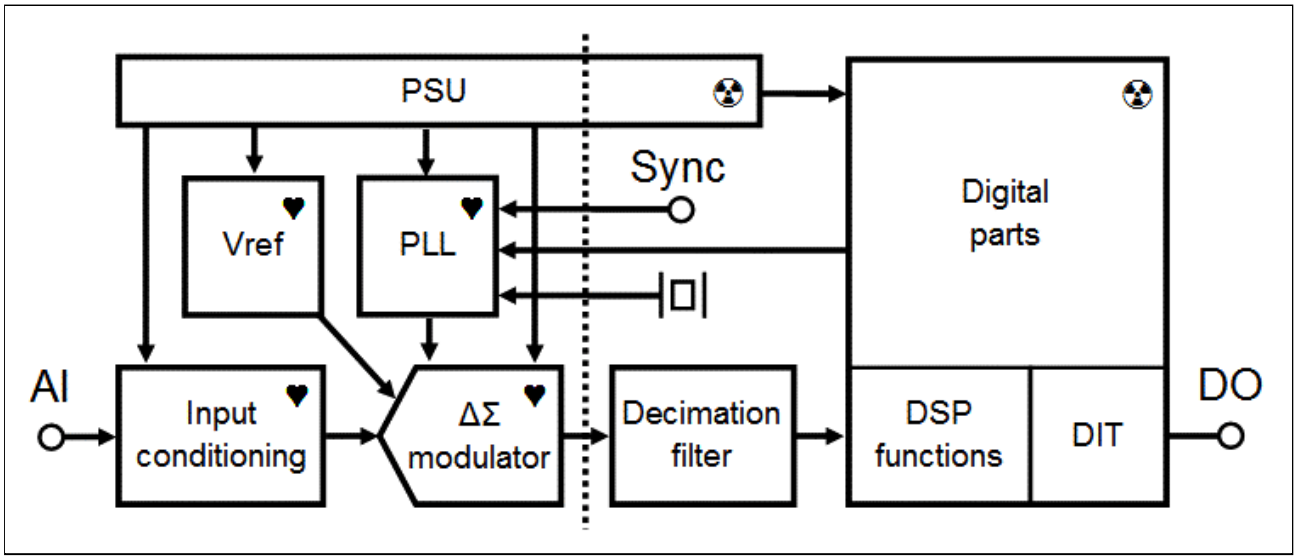

Figure 1: Typical ADC subsystem

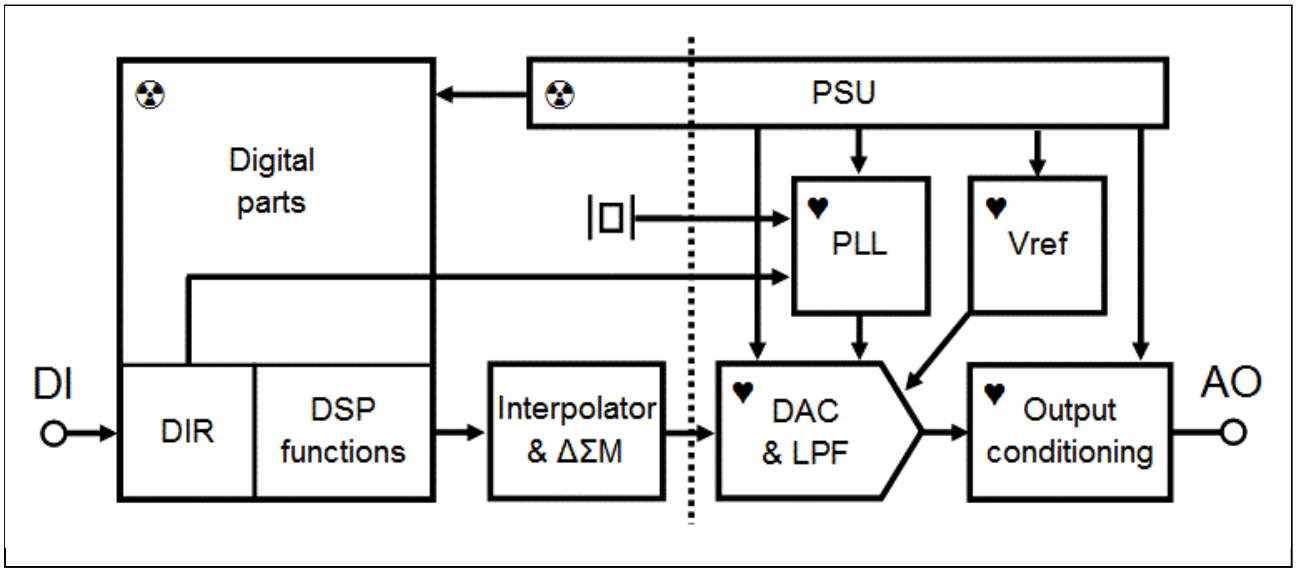

Figure 2: Typical DAC subsystem

Some thoughts about clock recovery

The conversion clock of a data converter is very critical to its linearity, since any variation in the regularity of the clock results in sampling jitter, which causes phase modulation of the converted signal [1]. This is most easily assessed by passing a high-amplitude, high-frequency signal through the converter and looking for phase modulation sidebands or skirts in the spectrum of the output.

Converter clocks are usually derived from a phase-locked-loop (PLL), rather than a local oscillator, since it is unusual for a device to be able to be system clock master all the time: it must be able to lock to an external reference or to its digital audio or computer interface. In order to lock to all references, the lock range may need to be wide (perhaps +/-1000ppm) and usually many different sample rates must be accommodated. See [2] for a chilling glimpse of the enormity of the problem.

The need to lock our converter clock to various wide-ranging references which maintaining low jitter has tended to be one of the most difficult challenges in converter design, since it embodies some tough tradeoffs. Although other applications (e.g. telecoms) had pretty much solved this problem before digital audio was thought of, we audio people just had to start from scratch and it's taken a couple of decades to reinvent a good solution.

The situation has become more challenging now that we need to lock to software-generated syncs and timestamps arriving over computer interfaces, since they can embody large amounts of jitter with uncontrolled spectrum, and may not come around as often as we'd like [3].

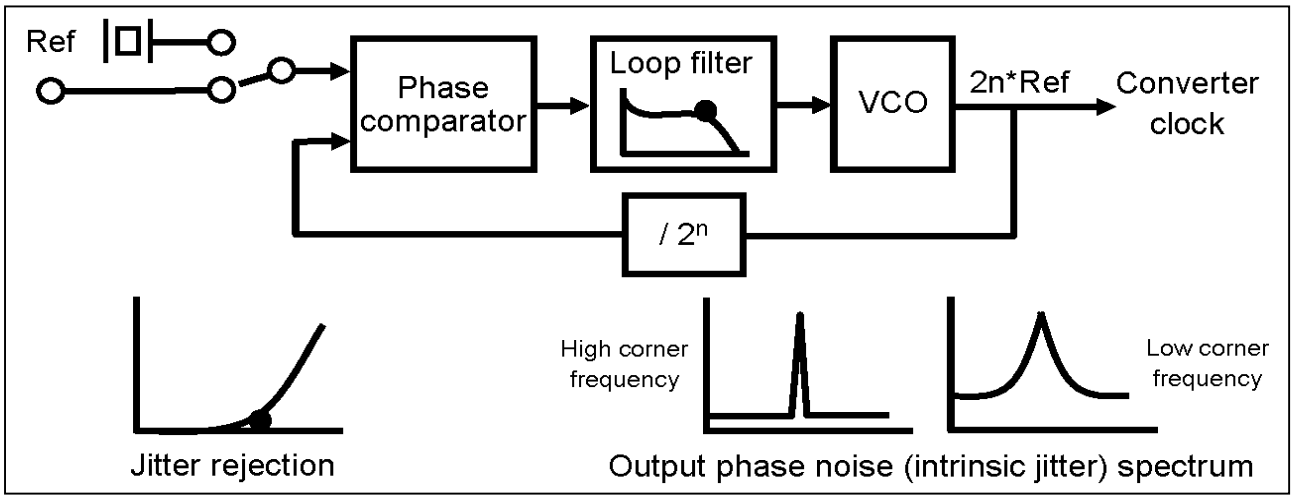

Figure 3: Basic analogue PLL

Figure 3 shows a conventional analogue PLL, such as might be used to lock a converter to a reference clock. A phase comparator continually compares the external reference with the regenerated version, and decides whether we should speed up or slow down our VCO to make them match in frequency and phase. But let's not be hasty: we need to respond in a leisurely fashion or else we will track any incoming jitter. So we smooth out the up/down requests with a low-pass loop filter before passing them to the control input of the VCO. So far, so good.

But the tough tradeoffs are mostly about choosing the right loop filter characteristics and the right sort of VCO. We need to reject incoming jitter ('jitter rejection' down to low frequencies in order not to be prey to audible sampling jitter [1]). That will also potentially allow us to accommodate a low comparison frequency (such as a reference comprising infrequent software time stamps, or a video sync which only lines up with an audio sample every few seconds [2]). Unfortunately a low loop filter corner frequency will make our PLL slow to lock up, but we can't have everything. Worse, though, it will prevent suppression of the phase noise of the VCO around the loop. So if we are to avoid unacceptable 'intrinsic jitter' we need to keep the loop filter corner frequency high, or else choose a type of VCO which has very low phase noise in the first place.

A good low-noise oscillator is a quartz VCXO. But because of their high-Q, they don't like to be pulled very far from their natural frequency and so may not have enough pull range for our requirements (perhaps

only +/-100ppm). On the other hand, a humble RC multivibrator VCO can have all the pull range we need, but is essentially untuned and so has large phase noise and is prey to all manner of interference. A few other VCO options exist in between these extremes.

To solve the pull-range problem with quartz, we could commission some special VCXOs made of a special material with a lower Q, e.g. Langasite (LGS) or Lithium Tantalate. These are quite expensive.

We could use a tuned-circuit (LC) VCO, which has much lower-Q than crystals, but at least they have Q, so they can be designed with much lower phase noise than a multivibrator. Their wide range can cover both n*44k1 and n*48k rates, and easily accommodate +/-1000ppm reference inaccuracy. Overall, not a bad choice; but if we're being picky, the intrinsic jitter isn't going to be top-class if we drop the corner frequency of the loop to where we'd like it.

In considering these tradeoffs, it is helpful to look at some real-world examples of applying analogue PLL technology to converter subsystems (Figure 4).

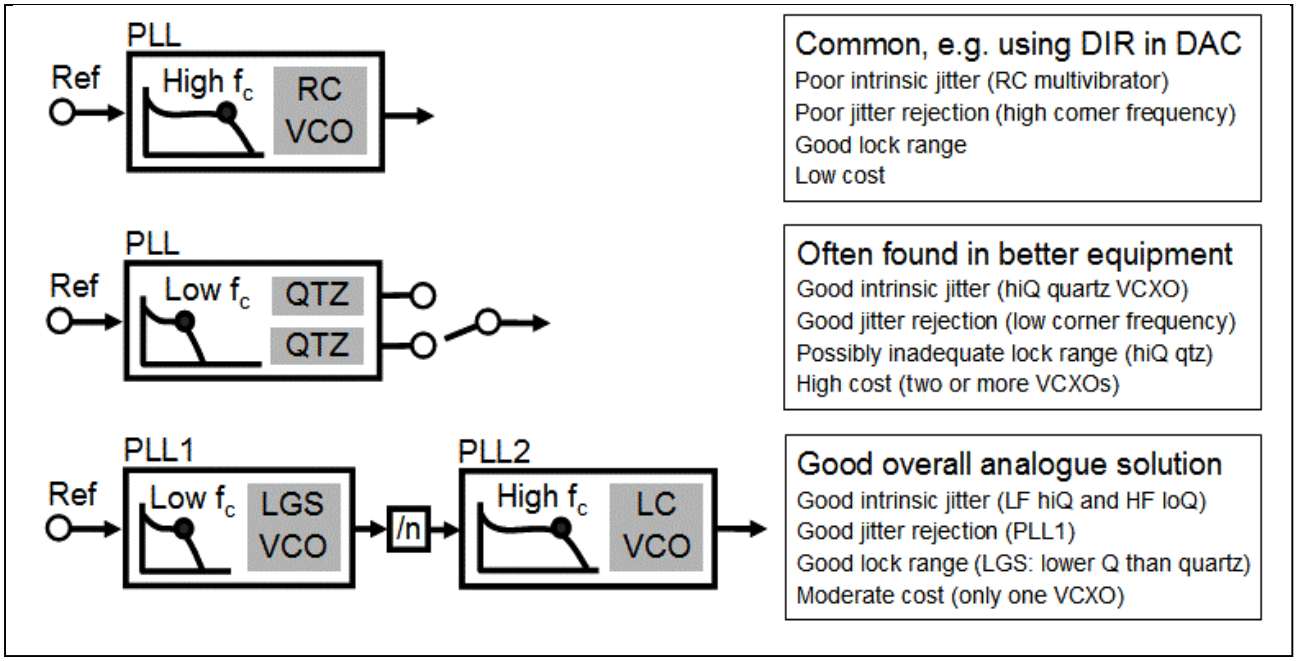

Figure 4: Some analogue PLL schemes

Top: using a basic PLL chip, or, for an AES3 or SPDIF DAC, using the PLL in the DIR. Problem: the VCO is low-Q with large phase noise (usually a RC multivibrator, which is vulnerable to all sorts of interference, especially power rail and ground noise), and the corner frequency of loop filter is high. Result: intrinsic jitter is high, jitter rejection at audible frequencies is poor.

Middle: using a purpose-designed PLL for converter clock recovery, with higher-Q VCO and a low loop-filter corner frequency. Result: intrinsic jitter and jitter rejection are good, but we have to carry the cost of two (or more) VCXOs and the pull range may be insufficient.

Bottom: using two cascaded PLLs; the first one with a LGS VCXO (for wide lock range) and a low corner frequency filter, which does our jitter rejecting; the second one perhaps with a tuned circuit VCO and a high loop-filter corner frequency (we don't need reference jitter rejection now, it's already gone) so as to cover all sample rate multiples. If we make the LGS VCXO frequency *n**48k and *m**44k1, we can provide a sample-rate reference for the second PLL with a simple programmable divider. Result: a pretty good solution; it has good intrinsic jitter and jitter rejection, and it works for all sample rate multiples of 44k1 and 48k, and has a wide lock range. But even with one VCXO it's still quite expensive, and it still may have questionable performance at low comparison rates.

An even better solution is to adapt the dual-loop architecture as a 'hybrid PLL', by implementing the first PLL in the digital domain, so that the trade-off of phase-noise vs. low corner frequency is broken: the loop filter and the 'VCO' are both entirely digital. The

VCO (now an 'NCO') is very jittery, since it is a varying integer division of a fixed master clock, but inclusion of a sigma-delta modulator in the loop means that its jitter can be confined to very high frequencies. It is therefore straightforward to cascade the NCO output into an analogue PLL with a very high corner-frequency, which can therefore use a cheap VCO without intrinsic jitter being a problem. Furthermore we can change the corner-frequency in software so as to achieve fast lock and then extreme LF jitter rejection. The hybrid PLL is very cheap, since it requires no resonator-based VCO.

Similar solutions are now available for audio use from a number of vendors (for example [4], whose topology is shown in Figure 5) – some of them are even built into data converters.

Finally, I should mention that with a bit of thinking outside the box there is another way to skin this cat. Modern low-cost sample-rate-converter chips (SRCs) are achieving performance which can arguably exceed that of the data converter itself. So you might elect to operate the conversion element at a fixed rate provided by a local crystal (thus eliminating sampling jitter) and to rate convert the converter input or output data. This approach can lead to other issues, and places the responsibility on the SRC to be able to achieve jitter rejection to the same standard as in the PLL model whilst also protecting the quality of your audio crown jewels.

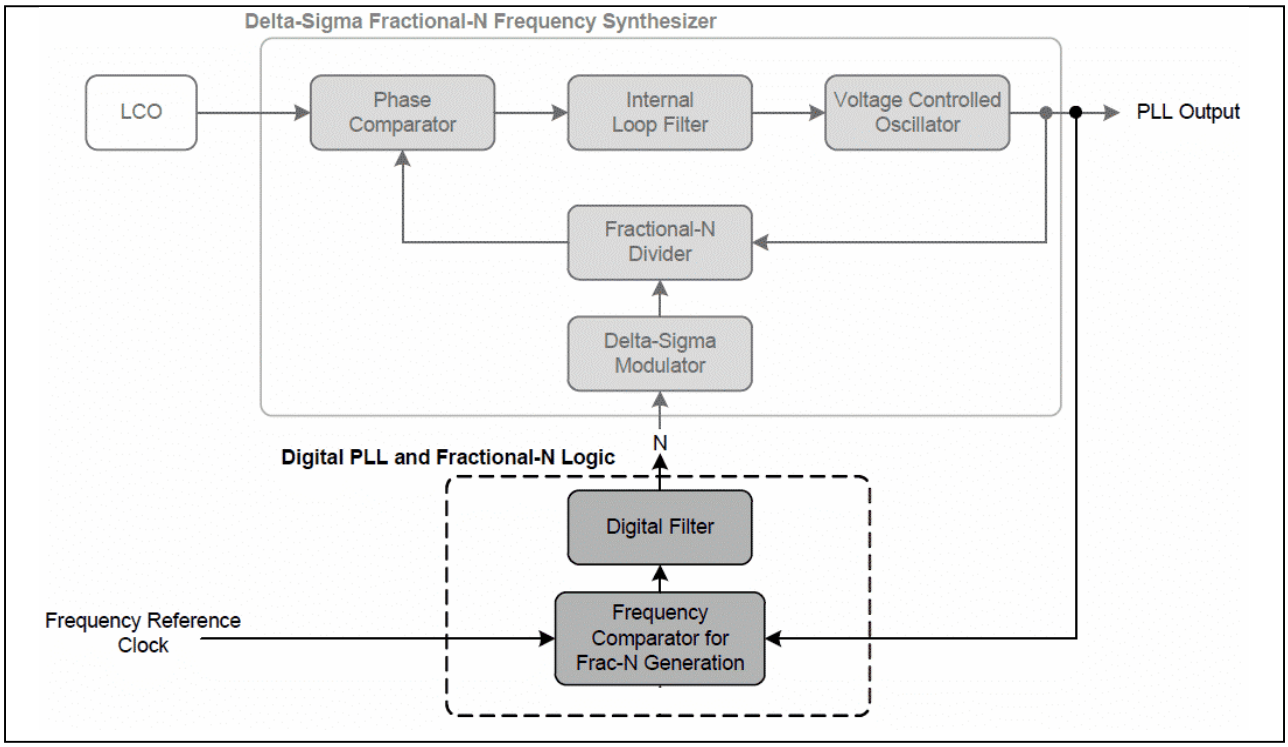

Figure 5: Hybrid PLL, from [4]

Some thoughts about power supplies

Linear and switching PSUs

High quality line-powered audio equipment has traditionally employed linear PSUs, but these have disadvantages of size/weight, heat and cost, and may need a manual line-voltage selector. But they do have the advantage (if properly designed) of not being a source of high-frequency interference into the analogue circuits. A switch-mode power supply (SMPS) eliminates these disadvantages, but must be carefully designed if it is not to become a major source of hostile switching noise. SMPSs have traditionally been feared by high-end audio designers.

A SMPS in any type of equipment must be designed to meet the appropriate EMC standards for conducted and radiated emissions. Nowadays this isn't very hard to do, with SMPS controller vendors providing application circuits designed to be compliant (usually only just, so as to save costs of filter components etc.). But an unfortunate fact of life for audio equipment designers is that the same sorts of misbehaviours which might cause a SMPS to be non-compliant can play hell with audio performance, even at much lower levels. Thus SMPS design for audio applications can be elaborate.

Drawbacks of low-cost SMPSs

Since converter systems often need a large number of power rails, and may benefit from isolation between the digital parts and the analogue, a 'flyback' architecture is a popular choice because it conveniently offers these benefits - which may also be useful in DC-powered situations, as discussed later. A wide variety of other 'simple' SMPS architectures can be used instead – the problem for the uninitiated is generally how to choose between them.

In a flyback converter, DC (either directly input or rectified from the AC line voltage) is switched though the primary of the 'flyback transformer' by a transistor under the control of a device which regulates the duty cycle of a train of switching pulses in order to keep the various secondary outputs regulated. The secondaries are rectified and filtered to provide the power rails.

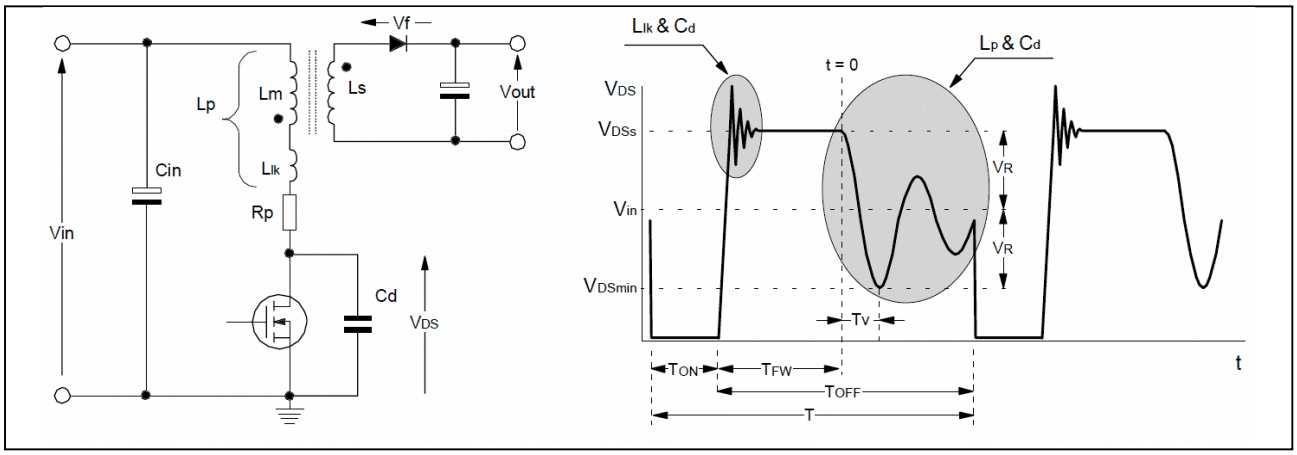

Figure 6: Flyback converter circuit and switching waveform, from [12]

Figure 6 shows a flyback converter, and its switching waveform. Note that the HF oscillation at the point where the switch is turned off is caused by the stray capacitance Cd across the switch, and the leakage inductance of the primary Llk, and is a largely-unavoidable source of hostile RF. Its amplitude can be controlled with snubbers to protect the switch, but this can even make the interference worse. The LF oscillation prior to turning on the switch is a consequence of the stray capacitance Cd and the primary inductance Lp, and is interrupted at a randon voltage by the switch-on, causing potential interference from the large switching voltage and current.

So the main problem with the basic flyback topology (and most other basic topologies) is that the transistor switches hard and randomly, causing high levels of radiated and conducted interference to invade the analogue audio parts. We are now on a slippery slope: there isn't much we can do to reduce the source of the problem (we can keep the switching loops as small as possible to reduce radiation, we can optimise the ground topology and make critical tracks fat to reduce ground noise; we can tame the edge times with snubbers, but only at the cost of reduced efficiency and increased heat). So we end up having to take disproportionate steps in the vulnerable parts to make sure that the audio remains clean. These can include screening cans, split grounds, galvanic isolation etc.

Another problem is the random frequency of the SMPS controller, which can produce interference at the beat frequency between itself and (for example) the audio sample rate, thus making it impossible to confine it to some inconspicuous part of the spectrum. A possible solution is to lock the switching frequency to some multiple of the sample rate in order to remove the beat frequency, but this can be problematic: the regulatory variation of duty-cycle can cause its own beat, and some types of SMPS controllers like to vary the switching frequency to effect regulation. You could even introduce a situation where a software bug or a wayward sample rate could collapse the power rails.

It's easy to see why we have been reluctant to use SMPSs in high-performance audio equipment, but in DC-powered situations we often have no choice, and SMPSs are small and cheap – so it would be good to find a way to tame them.

Resonant and quasi-resonant SMPSs

Ideally, to combine the benefits of a linear supply and a SMPS, we would like to find a way of passing a high-frequency sine wave through the transformer. An approximation to such a solution is a 'resonant' SMPS. In order for this to be achieved losslessly, it is usual to generate the sine wave by placing a resonant LC tank circuit in the primary – a neat trick is to use the primary inductance as the L part. Of course the stimulus for the waveform is still a hard-switching transistor under clever control, but the resonant circuit tunes the primary waveform to an approximation of a simple sinewave, with the switching happening at zero-voltage or zero-current moments, all of which leads to much less hostile switching noise.

On the other hand, resonant designs tend to be somewhat more costly and larger than flyback designs. A major drawback with many resonant architectures is that the LC tank circuit needs to be tailored to the switching frequency and DC input voltage in order to maintain resonant and zero-switching operation, which makes off-line universal input design problematic unless power-factor-correction (PFC) is incorporated.

A good compromise is a 'quasi-resonant' converter (QRC): the idea here is that since the problems in a simple flyback converter only happen at the moments of switching, it is only necessary to find a way of creating a resonant waveform at the switching points. This can be done quite simply by introducing primary resonance into an ordinary flyback topology, and making the controller clever about deciding when to switch. Figure 7 shows a zero-voltage-switching (or 'valley switching') QRC which is a low-cost way to cut SMPS emissions at source. The tank circuit (in this case created by simply adding a large Cd across the switch) slows the switch-off rise time, and the controller arranges the switch-on instant to coincide with a 'valley' in the LF oscillation; this inherently makes the cycle period variable with an attendant 'spread spectrum' effect which can improve interference and certainly increases EMC margins, as shown in Figure 8.

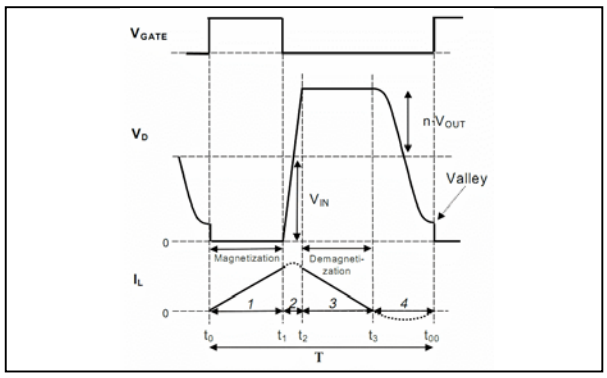

Figure 7: Switching waveform of flyback-based quasi-resonant converter, from [13]

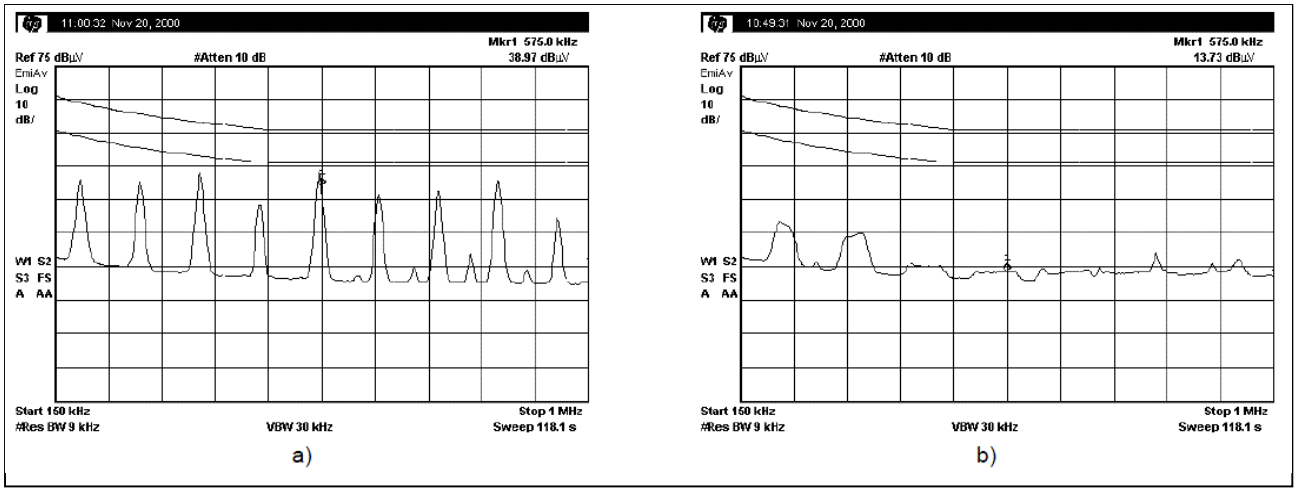

Figure 8: Conducted EMI for flyback (a) and quasi-resonant (b) designs, from [12]

SMPS topology selection is vital for audio performance

It is often difficult for the uninitiated to make the correct choice of SMPS architecture for audio, because most SMPS controller vendors organise their selection tables by power capability, on the assumption that the audio designer, like everyone else, will want to choose the cheapest solution for his power requirement. The recommended SMPS for low power applications like this (say <20W) will usually be a very basic, hard switching type because the power is low enough for its emissions to be retained below statutory EMC limits with minimal filtering, and for the losses resulting from its indiscriminate switching to be insufficient to set it on fire.

Resonant and quasi-resonant topologies tend to be recommended for higher power applications where the switching noise and losses HAVE to be controlled. But for audio, it's often a good idea to make a low-power implementation of a high-power topology in the interests of achieving minimum emissions – the extra cost can usually be saved in not having to armour plate the audio parts.

Other considerations

Cross-regulation is often a problem with multi-rail SMPS designs, since the regulatory mechanism of the controller can generally only operate on one rail. Varying load conditions on individual rails can cause the voltages of other rails to vary, an effect known as cross-regulation. In performance-critical applications it may be necessary to provide linear post-regulation on analogue power rails from the SMPS. If so, care should be taken that the linear regulators are adequately cooled. It should also be noted that linear regulators can usually only regulate over a limited frequency range, and the switching products from the SMPS can easily exceed this, resulting in their passing straight through the regulator. It is therefore recommended to use ferrite beads ahead of linear post-regulators.

Line powered linear and SMPS equipment usually draws current from the mains only during voltage peaks, which can cause the power line to be distorted with possible detriment to the performance of other sensitive audio equipment. It may be beneficial to design power-factor-correction (PFC) into you SMPS design in order to cause the least possible distortion to the power waveform (although you can be sure that all the other equipment around will probably be causing distortion anyway). Some PFC schemes allow tight control of the rectified voltage ahead of the SMPS, which can allow switching noise to be further reduced in resonant and quasi-resonant designs by ensuring zero-switching for any input voltage.

As well as applicable safety, EMC and disposal legislation, line powered equipment is already, or will become, subject to various territorial legislation for standby power consumption (where applicable) and operating efficiency (for example under the EU Ecodesign Directive and the voluntary US EnergyStar program). Although territorial legislations vary, maximum standby power consumption of 0.5W and minimum operating efficiency of about 80% are typical for a 20W device.

Some thoughts about the analogue signal path

I was wondering whether to bother with this section, since there is absolutely nothing here which will be news to generations of analogue designers. But then I remembered that analogue performance is nonetheless the limiting factor in many converter designs. Somewhere along the line we forgot some of this wisdom.

All of the buffer amplifiers, gain stages etc. between the outside world and the ADC, and between the DAC and the outside world are an obvious area where good design practice will pay dividends – it is no mean feat to maintain the performance of a flagship data converter through the analogue circuits, particularly if you have to incorporate significant functionality.

In general, use a ground plane for analogue circuits, and take advantage of the very small SMT packages which are now available. Lay things down instead of standing them up. These measures will reduce susceptibility to interference and crosstalk for free.

Of particular importance are the buffer circuits which drive by the input of the ADC or are driven by the output of the DAC. In general, the safest policy is to stick to the exact circuit topology and components recommended by the converter manufacturer. They will have spent a long time coaxing the best out of their device by tweaking the buffer. However, this is not always the case! With experience and care it is sometimes possible to exceed the 'application note' performance. On the other hand, if cost is important you can often scrimp a bit on opamp types – the manufacturer is usually more interested in squeezing the best out of his device than in your budget.

Sigma-delta ADC inputs often have a non-linear input characteristic and will produce aliasing components if subject to HF, so an ideal ADC buffer must achieve a good amount of HF rolloff (but without compromising in-band flatness) coupled with a low output impedance. So it's better to avoid the passive-pole between the output of the buffer and the converter input, and to use something like [5] instead.

Many high-quality DACs have current outputs, requiring an outboard current-to-voltage converter (IVC) circuit. Sigma-delta DACs often produce significant out-of-band noise, and the IVC must filter and/or cancel this in the first instance. Therefore the IVC must behave linearly up to very high frequencies (well above the audio band) if in-band linearity is to be maintained. If you are tempted to stray from the manufacturer's recommended IVC design, bear in mind the bandwidth requirement and note that the data converter output loading will probably have to be similar to the application circuit for optimum performance.

As for the rest of the analogue signal path: it is important to choose the right components and circuit topologies. The most straightforward way is to use opamps – and I'd say that nowadays this is a good policy except in a few very special situations such as mic/phono preamps and high-current outputs. It will pay to familiarise yourself with the noise models for opamp circuits (e.g. [6]) and to implement a spreadsheet to calculate noise levels for your particular circuit. Simpler, but less versatile, is to use opamp manufacturers noise reckoning tools (e.g. [7]). Best is to use a full-blown SPICE simulator (e.g. [8]).

You will find that your choice of opamps in each stage is generally restricted to relatively few which have the requisite voltage (and/or current) noise performance. The world is full of opamps, most of them no good for audio. On the other hand, beware the 'best audio opamp' syndrome. There is simply no opamp which will behave best in every audio stage. In each case, you need to select the right one for the job. Usually you can do this from data, by trading off particular requirements against cost, power etc. But don't be afraid to use trial and error in the end – although it needs some determination and good eyesight in this age of SMT. With some dexterity, you can persuade DIL sockets onto the SOIC sites in your prototype, and you're in tweaky heaven. In general, I like to use dual opamp sites – it is a good tradeoff of cost and choice, and allows tight layout in balanced circuits. Anyway, enough said about opamps – I don't need the death-threats, so I'm not going to recommend any.

Inverting opamp topologies are generally preferred to non-inverting, since the input terminals operate at a comforting virtual earth. The dynamic input-common-mode-voltage of non-inverting configurations may add distortion, particularly at high frequencies with some opamps, with others apparently not.

Consider adopting a fully balanced topology end-to-end (perhaps using a symmetrizer in the ADC case to achieve balance through the channel even with unbalanced inputs – but be sure to avoid mode-conversion). This not only reduces interference and crosstalk, but can also achieve higher performance since distortion mechanisms often tend to cancel. Most high-end data converters like to operate in a balanced mode anyway.

It is important to select a gain structure which maximises the dynamic range (SNR) and thus minimises the need for excessively low noise design, with its attendant cost. Even so, resistor values need to be low enough for their thermal noise to be out of the

picture, but not so low as to bring problems of over-dissipation or circuit loading.

Linearity of resistors and capacitors is also very important: resistors must be metal film or thin film types, not the usual chip resistors which are made of dirt and spit and change resistance according to voltage and the seasons. Capacitors must be low-k ceramic types (C0G/NP0) or low-loss plastic (e.g. polystyrene). Failure to observe this leads to non-linearity (distortion) in most circuit topologies. Electrolytics should be kept out of the signal path; there are other ways to ensure extended LF response.

PCB layout of the channel circuits is critical in order to minimise both interference and inter-channel crosstalk. Make sure that the opamp output and ground nodes of the stages are outermost and the opamp input nodes are innermost in your channel strip layout.

Pay attention to the bandwidth of your stages: it's not always better to go for a 'DC-to-light' approach; limiting the bandwidth at the top end reduces susceptibility to interference (and the presence of excessive HF, even if inaudible, does no good to the in-band signal); limiting the bottom end removes 'wandering DC' which can cause problems which mixing and switching, as well as the risk of unexpected overload. On the other hand, keep a healthy disrespect for the '-3dB 20-20k' approach. Extend the top and bottom end beyond the bare-essentials where you can, and strive to keep the area in between very flat – you will notice the difference.

Advice about the data converter reference voltage

Nearly all data converters have at least one accessible reference voltage (or current) pin. The input to the converter is multiplied by the reference to produce the output, and so what you do to the reference is just as important as the analogue signal path: any noise or interference on the reference will modulate the converter output. Internally-generated voltage references must be filtered with suitable capacitors placed very close to the pins – an assortment of low-value HF parts and larger electrolytic/tantalum types are usually required. It may be preferable in some situations to drive a voltage reference externally from a well-regulated and filtered source. Some converter devices require both high and low reference voltages to define their operating range, and some require separate references per channel. Whilst the actual reference voltages may sometimes be user-modified to some degree, converter performance is often optimised at a particular voltage so it's best to stay there.

A parting word: when you have distortion or noise problems, and you've looked everywhere else, don't forget the reference. It you have modulation issues (sidebands or noise skirts around the signal frequency): if it gets worse with increasing signal frequency, it's jitter; if it doesn't, it's probably the reference voltage. Check out [9] if you don't believe me.

Some thoughts about the digital parts

Processing

It goes without saying that any poor-quality digital processing in the signal path can undo all your good work in the analogue domain. Make sure that there is enough precision everywhere and that algorithms are beyond reproach. Particular culprits are often dithering/noise-shaping, and dynamics processing (remember to upsample).

Remember too those parts of the digital signal path which are often overlooked. For computer interfaces, the driver may let you down, or parts of the operating system that you thought you bypassed. It is important to be able to test your entire signal path, so make sure that your 'dual domain' audio test equipment can work in the 'computer domain', interfacing with your driver layers directly.

Interference

In the old days it was just about possible to keep all the digital parts of a converter system ticking over at some multiple of the sample rate, or a close relation – consider the simple case of an AES3 interfaced converter with no DSP, no microcontroller.

Nowadays, forget it. You are bound to have a box-load of computers in one form or another. DSPs, RISC processors, FPGAs - all running asynchronously, and probably a nasty computer interface buzzing away across the entire spectrum.

Although lower operating power, lower core voltages and ever smaller dies and packages contribute to reducing the hostile intent of digital electronics, there is little beyond this that the designer can do (apart from observing good EMC design at ports, and good power decoupling and filtering).

System considerations

Some fundamental decisions probably need to be made at the outset:

Power source: line-powered or DC-powered (wart or computer interface 'bus power')?
Habitat: can we put it in its own metal box, or must it cohabit with digital parts – or worse, must it go inside a host computer? Do I need/can I afford screening cans?
Construction: can I use a multi-layer PCB with groundplanes, small SMT etc.?

Isolation: can I galvanically isolate the analogue and digital domains using optical or magnetic isolators? This can be particularly beneficial in computer or computer-interface cases, where the digital ground can be very toxic.
EMC, efficiency: what are the statutory requirements? Can I afford to improve on them in the interests of audio performance?

The answers to these questions are probably dictated by project requirements and component budget, and you'll have to make the best of it. If audio performance is at all important, I'd say 'Do if you can afford it'.

ASSESSMENT

Objective assessment of converter systems, both during design and in general, is a hugely complex subject, and is the subject of many international standards such as [10] and [11]. However, it is useful in many situations to be able to make a simpler assessment, and for many this involves using listening tests instead.

The debate as to whether measuring or listening is best will rage forever. But to me, developing high-performance audio requires both methods. My own preference, as one who is devoted to transparency rather than any particular 'character' in audio equipment, is to use measurement to 'debug' the design and to worry about listening tests after that. With experience, I believe it is possible to get a feel for the relative importance of the various measured parameters in terms of audibility – up to a point. Although I know 'high-end' designers who don't possess any measurement equipment, it seems to me an impossible approach, since serious hard-to-identify issues will always mask delicate ones.

For now we'll stick to measured assessment, which is easy to apply interactively for 'debugging' during the design process.

Equipment

I would recommend equipping yourself with an audio analyzer which can stimulate and measure in the analogue, digital and 'computer' domains. It is a big advantage if the analyzer part can display a continuous high-resolution FFT whilst you tinker. If it's user-programmable, so much the better – set up a way of automatically hopping amongst the key measurements below. You'll want to be able to define all your key measurement parameters when you set it up. Make a lead to connect a pair of small probes to the analogue analyzer input, so that you can measure between all the stages.

For SMPS and computer-based devices particularly, you will also need a means of seeing the wideband spectrum. You'll need that for EMC pre-compliance anyway.

What to measure

I humbly offer Figure 9 as a guide to the most important audio performance parameters which apply to conversion systems. It isn't exhaustive. I suggest the 'transparency' issues in the upper part of the table as being worthy of particular attention. It's not that the other things aren't important, but they are not issues of transparency per se.

The most interesting parameters are dark-shaded. Again, it's not that the others aren't important: by definition, anything which degrades transparency becomes important if it's bad enough; but the dark parameters are the ones which are easier to mess up, and so are more important to keep a constant eye on – the others are less usually affected by design choices outside the data converter.

Transparency	**Linearity**	**Distortion**	**THD (vs freq, ampl)**
			IMD (close tone etc.)
		Frequency response	**Amplitude response**
			Phase response
		Nyquist	**Aliasing/imaging**
		Sampling jitter	**Intrinsic**
			Rejection
		Quantisation	**Low-level linearity**
			Dither, noise-shaping action
		Interference	**Modulation (RFI, powerline)**
	Additive	**Intrinsic noise**	**In band**
			Out of band
		Crosstalk	**Channel**
			Source etc.
		Interference	**RFI**
			Powerline
Other (not directly affecting channel transparency)		**Gain**	**Absolute & polarity**
			Interchannel
		Phase	**Absolute delay**
			Interchannel
		Analogue I/O	**Impedance**
			CMRR/balance
		Digital I/O	**Compliances**

Figure 9: Summary of converter performance parameters

REFERENCES

[1] Julian Dunn and Ian Dennis - `The Diagnosis and Solution of Jitter-Related Problems in Digital Audio' Preprint 3868, 96th AES Convention, Amsterdam, February 1994

[2] AES11-2009: 'Synchronisation of digital audio equipment in studio operations' (Revision of AES11-2003). www.aes.org/publications/standards/

[3] Julian Dunn – 'Sample clock jitter and real-time audio over the IEEE1394 high performance serial bus' Preprint 4920, 106th AES Convention, Munich, May 1999

[4] Cirrus Logic CS2000 datasheet http://www.cirrus.com/en/pubs/proDatasheet/CS2000-CP_F2.pdf

[5] AKM AK5394A datasheet, page 19 – Figure 9, 'Analog input buffer example 1' http://www.asahi-kasei.co.jp/akm/en/product/ak5394a/ak5394a_f03e.pdf

[6] TI Application Note SLYT094 'Calculating noise figure in op amps' http://focus.ti.com/lit/an/slyt094/slyt094.pdf

[7] TI NoiseCalc op amp noise spreadsheet http://focus.ti.com/docs/toolsw/folders/print/opamp-noisecalc.html

[8] TI "TINA-TI" free spice simulator http://focus.ti.com/docs/toolsw/folders/print/tina-ti.html

[9] Ian Dennis, Julian Dunn, Doug Carson – 'The Numerically-Identical CD Mystery: A Study in Perception versus Measurement', 1997 http://www.prismsound.com/m_r_downloads/cdinvest.pdf

[10] IEC 61606-3 –'Audio and audiovisual equipment - Digital audio parts - Basic measurement methods of audio characteristics - Part 3: Professional use'

[11] AES17-1998 (r2009): AES standard method for digital audio engineering - Measurement of digital audio equipment (Revision of AES17-1991). www.aes.org/publications/standards/

[12] ST Application Note AN1326 - 'L6565 QUASI-RESONANT CONTROLLER' http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/APPLICATION_NOTE/CD00004217.pdf

[13] NXP Application Note – '75W SMPS with TEA1507 QuasiResonant Flyback controller' http://www.nxp.com/documents/application_note/AN00047.pdf

[14] AES-12id-2006: 'AES Information Document for digital audio measurements - Jitter performance specifications'. www.aes.org/publications/standards/