

# Audible and Inaudible Problems with Digital Audio

## *Here are some musings on digital audio*

Digital audio was originally introduced to the world with the CD as "perfect sound, forever". Well, nothing is perfect, is it? However, let me say up front that in my opinion, digital audio can come very close to living up to that tag line. I'm not saying it always does (it doesn't).

So if it isn't "perfect sound, forever" what are its defects? And which of those are you actually likely to hear in the real world?

## Inaudible Defects and Enhancements

Below is my guess as to *real defects or enhancements beyond the CD standard* that you cannot normally hear under reasonable, practical, listening conditions. What do I mean by that? A room in a normal domestic setting with a high quality two channel sound system installed. The sound system volume adjusted to a high - but not uncomfortable - level. People in the room listening attentively and not engaged in any other activity.

Such a "normal listening room" is unlikely to have a noise level below 30dB SPL (as measured with an A-weighted noise meter). More sophisticated, frequency dependent, noise assessments are expressed as Noise Criterion (NC) levels, and a quiet room would be NC-30 (see [here](#) for useful information on this). Now, what follows will no doubt be considered contentious by some! And the only way to establish whether these issues are or are not audible would be, of course, double blind listening tests. For some of the issues noted here, this has been done. For others (e.g. whether missing dither can be heard) I am not aware that it has.

- **Dither (gone missing)** ... you won't notice if someone forgets to add this. Dither is an essential part of reducing the precision of samples. Any time you go from, say, 24 bits/sample to, say, 16 bits/sample, you should do so with added dither. Without dither, a very low level signal (which you will get as things fade to silence) will get truncated to square waves (albeit very low amplitude ones). They have a lot of high order harmonics and will sound terrible.

If a random signal, slightly smaller than 1 LSB of the reduced precision signal, is added to the full precision signal before it is truncated to lower precision, then low level signals in the full precision signal will be preserved in the duty cycle(s) of the least significant bit(s) of the reduced precision signal. This is dithering. Distortion (due to truncation) in the reduced precision signal disappears when dither is used, but at the cost of increased overall noise. An extra bit of cleverness called noise shaping can move the spectrum of that additional noise up to high frequencies where you cannot hear it as much. With dithering, signals that would be completely below the noise floor of the reduced precision signal (i.e. would not be there at all with straight truncation) will be preserved and can be recovered (e.g. heard - in principle).

So ... dither is obviously a good thing and has a bulletproof theoretical basis. But you will (I think) still *only hear the difference between dithered and undithered material if you turn the volume way up* as the music fades out ... turn it up to a level that would be deafening on the rest of the material. In a normal listening room at tolerable volumes, the benefits of

dither are probably almost always inaudible. For much more of this, with examples, see [this page](#) on Ethan Winer's excellent web site. Another indication on the general inaudibility of dither is that a software bug in a digital audio workstation product once caused dither to be silently (!) omitted even when it was turned on. This bug persisted for some while and none of the product's users noticed (see [this \(29th. April 2008\)](#) for more information).

Please don't misunderstand me - dither *should* be used every time a high precision sample is converted to a low precision sample. It is the only right way to do this. But under normal listening conditions, you won't hear the difference if it isn't done.

- **More than 20 bits precision** ... The fact is that the bottom 4 bits (at least) of the output of 24 bit A/D converters is noise. For example, if a 2V input gives a full scale (all-ones) output, the LSB corresponds to about 30nV. The bottom 4 bits is then covered by 480nV. The thermal noise generated by a 700 ohm resistance over a 20kHz bandwidth at 20C comes to that! Some seriously careful, low impedance design would be needed to make those bottom 4 bits truly meaningful. In practice, of course, it can't really be done.

Since 24 bit converters are the norm now, there is nothing to be lost by using them. At least it isn't going to be necessary to add an *analog dither voltage* to the signal, as was advisable with 16 bit converters! That thermal noise will do the job nicely! Other problems with A/D converters get smaller as the number of bits goes up, so using a 24 bit converter is certainly not a bad thing. In reality, though, 18 bits of data is probably the best you can hope for before noise takes over.

You won't hear any difference between 20 bit and 24 bit precision. Under normal listening conditions, you won't hear the difference between 16 bit and 20 bit either.

- **192kHz sampling rate** ... Higher sampling rates have long been popular as a cure for (supposed) "digital ills". The objective justification for going beyond 44.1kHz begins with the need for anti-aliasing filters in the analog signal path to remove signal components above half the sampling frequency. Everything above 22.05kHz must be very heavily attenuated to avoid potentially very nasty sounds.

Designing an analog filter that has no attenuation at 20kHz and, say -100dB attenuation at 22.05kHz and doesn't have any nasty phase shifts in the audio band (and is otherwise generally well behaved) is very difficult. The components required are also expensive. Moving to 48kHz gives us 4kHz rather than 2kHz over which to go from zero to complete attenuation which helps a little. You can see why such circuits are called "brick wall" filters, though! Because some compromises pretty much have to be made in such brick wall filters, these were a prime target for those who thought they heard something wrong with digital audio as it was introduced in the 1980's. In this case, they may sometimes have had a valid point! Moving to 96kHz sampling rate gives us 76kHz in which get sufficient attenuation. This is much easier to design and implement.

So - the higher the sampling rate, the better? Well, actually, no. There are a number of issues with increasing the rate, and 192kHz looks to be a step too far. For starters:

- The good thing about higher sampling rates is that it allows "gentler" anti-aliasing filters to be used. But, because they are gentler, they will let more ultrasonic content

pass. You cannot hear that content as such, but it can cause audible effects through inter-modulation in various parts of the audio chain. Whether there are any audible effects depends on the details of the whole system, but it is certainly a possibility.

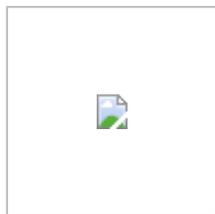
- Thermal noise increases with increased bandwidth. Using higher sampling rates and gentler low pass filters will inevitably increase the unavoidable noise floor of the system. Audibly? Probably not under realistic conditions.
- 192kHz sampling requires twice the storage of 96kHz and four times the storage of 48kHz. This is stating the obvious and loss-less compression may result in considerable savings over these ratios, but, even with disk space being almost free, it might be something to think about.

Since the ear can only hear up to 20kHz (OK - maybe a little more in very young people, with a following wind), it would seem as if anything beyond 96kHz sampling is, at best, pointless and, at worst, actually harmful. There is also another factor - [oversampling converters](#) allow the use of the lower sampling rates with gentle low pass anti-aliasing filters anyway! In the case of D/A converters, the digital signal is interpolated up to a higher sampling rate (digitally - of course!). This interpolated signal will not contain any new high frequency components (at least, ideally), and its aliases are spaced apart by the higher sampling frequency of the interpolated signal, allowing the analog anti-aliasing filter to be as gentle as with 192kHz sampling - or even more so. With this technology - now universally employed for many years - even 48kHz and 44.1kHz sampling rates can avoid the "brick wall" filter problems that were suspected of causing offence in the 1980's.

A very thoughtful discussion of potential problems with 192kHz sampling rate audio can be found [here](#) (very relevant to the "iPod age") and another [here](#). With oversampling converters, I doubt that anything more than the 48kHz sampling rate is needed. You won't hear any difference between the different sampling rates.

- **Jitter** ... This is another oft cited cause of unpleasant sounding digital audio. It is assumed in sampling theory that the samples are taken at exactly equal intervals of time. In the real physical world, however, nothing is "exact" to infinite precision. The sampling clock will vary slightly in frequency. This means that, assuming the A/D and D/A converters are otherwise perfect, you will be getting the "right" value, but at the "wrong" time! After digitization, you will get the wrong digitized value if the time interval is shorter or longer enough to move the value by 1 least significant bit compared to what it would have been if the interval had been exactly what it should be. The worst case for this is a sine wave at the maximum anti-alias pass-band frequency (conventionally: 20kHz).

How big must the sampling clock jitter be to get the wrong digitized value? The answer is:



seconds

where  $f$  is the 20kHz sine wave frequency and  $n$  is the number of bits in the converter. It turns out that the jitter needs to be very small: 121ps for 16 bits, 7.6ps for 20 bits and (wait for it!) 0.5ps for 24 bits! Yes - "ps" means "picoseconds"! I was amazed many years ago when I came across this - who would have thought that such *short* time intervals

could be relevant to something as *slow* as digital audio? Of course, a 20kHz sine wave is the worst case ... but even so.

The standard technology for generating stable timing signals is the quartz oscillator. The good news is that the **basic jitter** of these is very good: of the order of 1ps and some are specified to have less than 0.2ps jitter. So there is some hope of meeting these extreme stability requirements - at least in theory - although in reality very careful design and implementation would be needed to get anything like such levels.

The bad news is: a quartz clock can only be used *directly* for the A/D converter. The D/A converter must follow (with some leeway) the incoming data stream clock. This is why the digital *interface jitter* can have some effect on the *sampling jitter* we have been talking about so far. It turns out that in common interfaces (e.g. S/PDIF) the jitter of the clock recovered from the incoming data stream has the nasty characteristic of depending on the data and is generally higher than is desirable. The way around this is to use the recovered clock as the reference for a *phase locked loop* that follows the incoming clock but filters out the jitter with a low pass filter in its internals. A typical S/PDIF interface (Wolfson 8805) specifies 50ps jitter.

This doesn't sound that great, does it? But it *is* possible to use a quartz clock to drive the D/A using an *asynchronous sample rate converter*. The input to the sample rate converter is clocked by the PLL regenerated clock, but the output is clocked by the quartz oscillator. An example of such a device is the Analog Devices AD1896. The use of such devices should (at least in principle, everything else being equal) get the D/A jitter down to the same level as the A/D jitter. Phew! So ... with some difficulty, we can have sufficiently low jitter to make (maybe) 20 bit conversion of the worst case signal meaningful (on a very good day). It is probably more realistic to expect a few tens of picoseconds of jitter, though. (Interestingly, both thermal noise and jitter considerations converge on a similar figure for the number of truly significant bits you are likely to have: around 18).

What level of jitter do careful experiments tell us we can actually hear? A careful **study** found that none of their test subjects could detect random jitter of 250ns or less. That is *several thousand times higher* than the levels we can hope to approach and that *will* start to introduce *real, measurable, errors* with the worst case signal. Other studies have come to similar conclusions.

What sort of jitter do digital sources you can go out and buy have? Well, the validity of a lot of the data I can find trawling the Internet isn't that clear, but it looks like a humble Squeezebox may deliver around 60ps, with "a cheap CD player" coming in at 270ps and "the worst ever measured" at 1100ps. So - the "worst ever measured" is about 1ns ... 250 times *lower* than seems to be audible.

Wow! It seems very unlikely that you will hear any difference between sources due to their jitter performance.

## Audible Defects and Deficiencies

So, what problems *can* arise with digital audio? Here are some thoughts on that ...

- **Lossy compression artifacts** ... Once we leave the safe world of uncompressed or

losslessly compressed audio, pretty much anything is possible! Albeit very unlikely, if used sensibly ...

Lossy compression algorithms rely on the psycho-acoustic phenomenon of *auditory masking*. Extensive studies have revealed how the ear/mind cannot perceive low level tones at certain frequencies in the simultaneous presence of other higher level tones at other frequencies. If the ear/mind cannot hear some components of a signal, they can be thrown away and no one will notice. The software that decides what to discard and then compresses (losslessly) what remains is called a *coder* and it relies on a model of what the ear/mind cannot hear. The compressed data stream is passed to a *decoder* which reconstructs the original signal with the inaudible bits (according to the ear/mind model) missing. Often, the coder and decoder software are both called a *codec* which I find a bit confusing, since the coder and decoder are not the same.

Lossy coding will work so long as (a) the ear/mind model is sufficiently accurate to predict what can be safely thrown away, and (b) the user does not insist on reducing the coder output data rate to a point where things that should not be thrown away are discarded anyway.

Lossy audio coding was first proposed in 1979 and first implemented in the late 1980's. It became popular starting in the early 1990's and was primarily standardized under the auspices of the ISO Motion Pictures Experts Group. A number of standards emerged, in fact. The most important were MP2 (more precisely MPEG-1 (or -2) Audio Layer 2) from 1993 and MP3 (MPEG Audio Layer 3) also from 1993 (although it actually followed on from work on MP2 which began as early as 1987). MP2 is important for DAB (Digital Audio Broadcasting), while MP3 has taken over the world (or so it sometimes seems)!

Given a sufficiently high allowed compressed bit rate, both MP2 and MP3 can produce excellent quality. The goal for MP3 was to do as well at 128k bits/sec as MP2 did at 192k bits/sec. Implementations of codecs have improved over the years too.

It seems generally accepted (although probably not universally) that at 320k bits/sec, on pretty much any material, MP3 is indistinguishable from 16 bit, 44.1kHz uncompressed CD quality signals. I haven't listened to too much MP3 material in "normal listening conditions", but I haven't heard anything untoward at the bit rates I have used. These range from 94k bits/sec (mono, speech) to 244k bits/sec and result from using Lame to transcode from FLAC format using `--preset standard` which results in variable bit rate output. MP3 seems to be very good indeed.

When I listen to MP3, though, it is usually under far from ideal listening conditions (e.g. on an airplane or while mowing the lawn) ... I suspect that the vast majority of MP3 use is of this kind. More on this later ...

In my experience, there can be real issues with MP2 based DAB:

- At 192k bits/sec, I would say it is perhaps marginally superior to a good FM tuner.
- At 160k bits/sec, it is close ... perhaps slightly worse.
- At 128k bits/sec, it is very definitely inferior. Seriously so on some material. In fact, sometimes it is just *awful*.

I have no idea what DAB is like at 256k bit/sec, because nothing is ever broadcast at that rate in the UK. I believe it is in mainland Europe. Apart from BBC Radio 3, pretty much everything seems stuck with 128k bits/sec here. Fortunately, I pretty much only listen to Radio 3! But really ... there is no way that FM should be turned off in favour of DAB. With portable receivers it is far too fragile, and rendering probably hundreds of millions of perfectly good FM receivers useless would be crazy.

- **Overworked error correction** ... Digital audio typically uses several layers of error correction whenever this is appropriate. As with all digital data processing, there are correctable and uncorrectable errors. Correctable errors will be corrected (!) and data will be recovered as if no error had occurred. Correctable errors may indicate impending problems with media, but they will have zero impact on the data - i.e. the audio. Uncorrectable errors *will* affect the audio. In the computer world, an uncorrectable error normally stops whatever processing is taking place. You don't want to continue with uncorrectable errors in your bank balance, for example! With digital audio, though, attempts are often made to patch things up as best as possible and continue. After all, what else are you going to do? But the results can be truly awful when this takes place.

Once again, DAB sticks out like a sore thumb as the place you are most likely to encounter this problem. When you hear something like the programme you are listening to being played under water, that is the result of patching up uncorrectable errors.

- **The fun world of computer audio** ... Where to begin? Computers are much the same with audio as with everything else they are used for: so useful they are virtually essential and at the same time incredibly frustrating and annoying - at least, at times. The audio specific issues are:
  - Will a soundcard work today, or won't it? This applies mainly to USB soundcards on Windows. I've had a couple of different models, used with more than one computer, and whether the device works reliably seems to depend on the phase of the moon. With my current USB soundcard (E-MU 0202) I finally gave up and bought a cheap secondhand laptop, loaded it with Ubuntu and haven't any issues since. But that may be just good luck. (I can be sure from decades of experience that all the current mainstream operating systems - Windows, MacOSX and various flavours of Linux - are just as good and bad as each other ... just in slightly different ways).
  - Which brings us to manufacturer support (for soundcards) ... An additional reason I gave up on Windows with the E-MU 0202 was that Creative Labs announced they would not be updating the driver for any future versions of Windows. I wonder how many perfectly good devices will end up being discarded because the makers drop software support? It makes commercial sense, of course, and it is reasonable to *eventually* drop support but some are very quick off the mark! Open source is the one sure way around that.
  - The various control panels for sound devices on Windows ... with all those minuscule sliders. One twitch on the mouse and maybe we will be doomed? How might the different panels interact with one another? Might I suddenly get some system generated beep or farp in the middle of something I'm recording? What do things like "Speaker setup" settings actually do? The list is endless ...
  - On Linux, how many different (if related) audio systems do we need? OSS, ALSA, Pulse Audio, JACK ... Another seemingly endless list ... More far from clear potential interactions ...

- More generally - Can the operating system (whichever one it is) get out of the way and let applications handle audio data correctly? The answer is "yes", but being confident that everything is working, reliably, and as it should ... that is not an easy place to get to!
- **Poor implementation** ... Both the hardware and software components of a digital audio system need to be implemented carefully with the best (or, at least, most appropriate) algorithms. It is one thing to realize that you can't hear missing dither under normal conditions, but quite another to find the designer just couldn't be bothered doing it! Things like that should be done! Some of them - interpolation and filtering algorithms come to mind - could certainly cause audible effects if done badly. Good power supplies and careful PCB layout with ground planes are essential to keep digital and analog signals apart and to keep jitter low (even if you won't hear it anyway).
- **Longevity** ... "Perfect sound, forever" ... Forever? Nothing is forever - certainly not digital data (which is all that stored digital audio is). There are two problems here and they are both really serious:
  - The media the data is stored on decays in some way so it can no longer be read. The original, pressed, CD disks actually will last a very long time, in all likelihood. A century or more should certainly be possible. Recordable CDs, on the other hand, rely on chemical dyes and may not be very stable at all. Things like leaving them in the sun will kill them very quickly. Hard disks (which have become the dominant storage medium for everything) are the least reliable of computer components. They are also so large that the only thing you can copy them to is ... another hard disk! Flash memory may be significantly more reliable than hard disks ... time will tell. Flash memory will slowly replace magnetic hard disks over the next several years.
  - The second issue is that whole classes of media become obsolete and the equipment to read them is no longer manufactured. This has already happened to many digital formats: Digital Compact Cassette (DCC), Digital Audio Stationary Head (DASH), Digital Audio Tape (DAT), PCM audio on U-Matic video tape ... all of these and more are now on the "seriously endangered" list. Several companies have set themselves up to preserve audio - digital and otherwise. One of them has a very interesting [website](#), full of useful information.
  - Actually, there may be a third issue. Software can become obsolete and unavailable for various reasons and this could perhaps have similar effects to media and hardware obsolescence. For straightforward uncompressed PCM audio and common formats such as MP3 this can't really be a problem. But if we extend audio formats to include project data files for digital audio workstations (DAWs), this is definitely a real issue. Many of these will be in proprietary formats, often with very complex internal organizations. If the companies who wrote this software go out of business, or even fail to provide a way of moving old projects to new versions of the software (think version 1 to version 10 here, rather than consecutive versions) huge amounts of work may go down the tubes with no hope of recovery. Again, open source (or at least open data formats) is the only solution to this.

There is only one answer to the problem of preserving digital data - copy it ... to many devices (ideally of different sorts) and keep the copies up to date! The more copies you have the more likely your data (i.e. music) is to survive. The good thing is that it is easy to

copy digital data, and the copies are identical (barring failures) to the original.

Of course, analog media have problems with longevity too. You cannot copy them without generation loss, which is a very serious issue to start with. And even basically very robust recording systems (such as analog tape) have had major issues with flawed media (sticky shed syndrome, for example). However, I think it may be fair to say that leaving a tape (and certainly a vinyl record) on a shelf for 50 years is likely to be a lot safer than treating a hard disk in a similar fashion. Digital data must be "kept alive" by frequent copying to have any confidence it will be available when you need it.

- **Intent** ... This is an odd one. But digital technology can be used in two quite distinct ways, and, given limited bandwidth and/or storage capacity somewhere in the system, a choice must be made as to which way to go. The two paths are: *Quality* or *Quantity*. When it was first introduced, the emphasis was on quality. But, over the years, that has changed quite dramatically to quantity. This is why DAB is broadcast at inadequate bit rates in the U.K. This is why MP3 files are often compressed too much - people want as much music as possible on their MP3 players. Do people not care about quality then? Well, I'm sure they do, but it comes further down the list than choice and convenience, I think. There is a good reason for this: most people listen to music most of the time in environments in which you really can't appreciate high quality anyway. Typically in the car or otherwise on the move. Very few people spend significant time in a quiet listening room appreciating beautifully reproduced music. Some do, sometimes - but not very many. That is just the way we live.

## Could Digital Audio be Fundamentally Flawed?

I am tempted to just say "NO!" and move on ...

However, there is a tiny but very vocal group of people who seem convinced that it is! Now, let's be clear: the vast majority of the public either couldn't care less or are quite certain that "digital = good". The small fraction of the population who disagree have had a very strong influence (to say the least!) on the specialist "hi-fi" press for a couple of decades though. I'm not going to attack the "subjectivist" school of thought at length here. That has been done very thoroughly, and much better that I could do it, by others. If people think they can achieve audio nirvana by extracting the last possible drop of quality from a scratched piece of plastic, good luck to them. After all, I am quite enamoured of rust stuck to some plastic ribbon, so I can hardly complain! So long as their views are honestly held, everyone is entitled to their own opinion on these matters. (Let's face it - it is a pretty harmless debate as debates go!).

The only question I will ask here is: Is there any reason to think that digital audio, as it is currently standardized, might be missing out on something critical to the human auditory experience? And, as a secondary question, could analog audio possibly do this "something critical" better?

We believe we understand the ear/mind very well indeed. Contrary to what is sometimes claimed, what can be heard by humans is not a deep mystery yet to be plumbed. The parameters selected for digital audio were chosen based on the known qualities of human hearing. In particular, the ear/mind is believed to be a band limited system with a well defined

frequency range, outside of which it simply does not respond at all. This justifies using a strictly band limited representation - which is what you get with digital audio. Is there any reason to think the ear/mind is not strictly band limited (to the ~20Hz to ~20kHz range)?

As far as I know, there is only one scientific study which could cast any doubt at all on this. An effect which the researchers call the "hypersonic effect" has been investigated at some length by a Japanese group of neurologists. See [Note 5](#) for one reference. It seems that measurements of brain activity indicate that there is a different brain response to sounds when those sounds include ultrasonic frequencies, and the response is clearly distinct from when they don't.

I'm not sure how highly regarded this work is, as it is far from my fields of knowledge. It certainly seems to be well thought out research published in respected, peer reviewed journals, and there seems no reason to suspect it in any way.

It isn't clear what implications this result might have for audio recording and reproduction - if any. It seems that further experiments have indicated that the ear is not responsible for the "hypersonic effect", and it isn't clear what might be. It could be that the body can sense ultrasound - as it can sense infrasonic vibrations below 20Hz. Or maybe not ... these results are pretty extraordinary and need to be scrutinized carefully.

As far as objective reasons to doubt the adequacy of digital audio for representing everything that humans can possibly hear (at least with so called "high resolution" formats), that is all I have.

So to the secondary question: could analog audio be "better" in some critical respect? Well, it is not so strictly band limited as is digital audio. Analog tape recordings often preserve signals well beyond 20kHz. But, again, that is all I have. Analog audio can sound wonderful, but I have no reason to believe it has anything apart from this that could make it "better than digital". And I don't see how signals you can't hear could really make things "better" anyway (neglecting this "hypersonic effect" research).

In truth, the weakest links *by far* in the electrical audio signal chain are the transducers at the beginning and at the end: microphones and loudspeakers - but especially loudspeakers! That is where to spend effort in making real improvements, in my opinion. Although a lot of loudspeakers already sound pretty darned good too! Objectively though, loudspeakers have much worse distortion performance than any other item (even analog tape machines and vinyl records). It is even possible that the "hypersonic effect" may be due to *intermodulation distortion* in the speakers reproducing ultrasonic signals generating sound in the normal hearing frequency range which the subjects hear in the normal way. A careful reading of the research papers would be needed to see if that possibility had been eliminated.

Above and beyond all components in the electrical audio signal path, however, the *acoustics of the listening room and the recording venue* have a much greater effect on the sound you hear. Much more than even the transducers. The pressure field in a room is, in fact, rather a "mess". That the ear/mind can make so much sense of this "mess" is quite an achievement.

## So Why is Domestic Audio Not Completely Digital?

What I mean by that is: Why aren't all the signals between components digital signals? Why isn't all source selection and processing done in the digital domain? Why isn't the only analog path in a hi-fi setup the one between a per-driver power amplifier and the driver itself inside a loudspeaker? After all, all modern sources are digital (at least, if we accept DAB as a suitable replacement for FM radio)?

It is possible to buy a system that is entirely digital in this sense. It has been possible since the mid-1990's, in fact. I have owned an example myself for more than a decade. But this approach is not at all common. In fact, only one manufacturer I know of seems to have gone down this route and stuck with it. By now, I had expected it to be pretty much universal. Given the nature of digital audio and the performance of low cost processors, there don't seem to be any real obstacles to doing this economically (OK - one power amplifier per driver might not be the cheapest approach, but with Class D amplifiers and other developments even that has become a lot more affordable in recent years).

I don't really know why sources are still "digital islands" in essentially analog systems. But that is definitely the way it is. Maybe it is because digitally interconnected devices from different vendors do not "play well together" for various reasons, and people want the flexibility to mix and match components from different suppliers. But I am surprised that any such obstacles have not been overcome by now and that there are still so few "all digital" options today.

I think I know why "all digital" is not popular in DIY audio. It is just too hard for various reasons: lack of suitable easy to use building blocks, and the need for software and well as hardware skills, for example.

---

Note 1: Oversampling converters are actually implemented using a technique called delta-sigma modulation. This is quite complicated. but a really great description of how it works is available at Uwe Beis's site [here](#).

---

Note 2: See:

- "Detection Threshold for Distortions Due to Jitter on Digital Audio", Kaoru Ashihara, Shogo Kiryu, Nobuo Koizumi, Akira Nishimura, Juro Ohga, Masaki Sawaguchi and Shokichiro Yoshikawa, *Acoust. Sci & Tech.* Vol. 26, 1, 2005

Also:

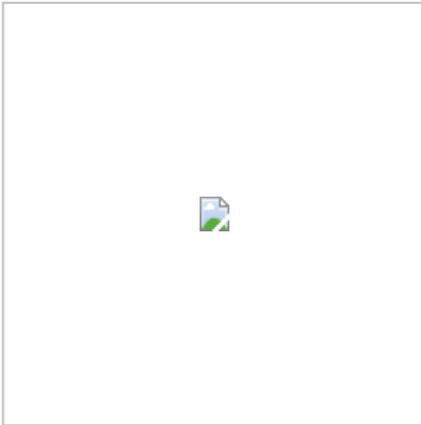
- "Theoretical and Audible Effects of Jitter on Digital Audio Quality", Benjamin, E. and Gannon, B., Preprint of the 105th AES Convention, #4826, 1998.

---

Note 3: How is jitter measured? The measurements here are given as r.m.s. time variations. Jitter can also be seen in the frequency domain - with a spectrum analyzer. In that context, it is usually called **phase noise** and it is far easier to measure the phase noise with a spectrum analyzer than to try to make time interval measurements directly.

Given a perfectly stable signal (say at 1kHz), the effect of jitter is to create sidebands at the jitter frequency. That is, jitter *frequency modulates* the signal (or music).

Jitter can also be seen as a *noise source*. The signal-to-noise ratio due to an r.m.s jitter of  $T$  seconds is:



For a 1kHz signal with a 60ps jitter, the SNR is 128dB falling to 102dB at 20kHz. For a 1.1ns jitter at 20kHz, we get 77dB SNR.

With frequency domain measurements, the power spectral density of the phase noise can be converted to an equivalent time domain r.m.s. jitter.

---

Note 4: It seems that not all errors are detected ... This is perhaps the dirtiest of the digital age's dirty little secrets. This *silent data corruption* is rare, but it does happen and rather more often than many would like to admit. With very large data sets, this is a very worrying issue. For scary details, see:

- "Keeping Bits Safe: How Hard Can it be?", Rosenthal, D. S. H, ACM Queue, October 2010.

---

Note 5: See:

- "The role of biological system other than auditory air-conduction in the emergence of the hypersonic effect", Tsutomu Oohashia, Norie Kawaia, Emi Nishinac, Manabu Honda, Reiko Yagia, Satoshi Nakamura, Masako Morimotoe, Tadao Maekawah, Yoshiharu Yonekurai, Hiroshi Shibasaki.

[Go home ...](#)