




Tech Note

GIBBS TRANSIENT OPTIMISED DIGITAL FILTER; TRANSIENTS LIKE *REAL* MUSIC.



INTRODUCING THE IFI GTO™ DIGITAL FILTER

ALL DIGITAL FILTERS FOR AUDIO ARE WRONG.
ALL OF THEM, INCLUDING THE 'NO FILTER' OPTION.
THIS IS WHY WE NEED YET ANOTHER FILTER!

All digital filters (including no filter) differ in *how* they are wrong and how this influences objective measured performance as well as subjective listening performance with music and indeed, specific types of music. All digital filters add specific distortion signatures in either time vs. amplitude domain or frequency vs. amplitude domain. These distortions become all the more relevant the lower the sample rate. So, the most abundant digital music source - CD quality - is most impacted with greater possible audible consequences than High-Res content.

Wherever there is a difference, there is also a preference. Subjective listening preference may be informed by a range of factors including a learned or acquired response to recorded sound (e.g. what sounds 'right' or 'hifi' is not what sounds natural in comparison to a live performance), including direct referencing acoustic music performances.



However, with sufficient data from extensive listening tests and some inductive thinking, one should be able to propose and implement a digital filter that offers substantial improvements in removing ultrasonic noise over the 'no filter' (non-oversampling) case while avoiding as much as possible erring too far in

the other direction with excessive and audible ringing.

So here it is - the ever so musical iFi GTO™ Digital Filter in the Pro iDSD which is the first ever seen in any DAC. In due course, technological hurdles permitting, we will try to implement it as a firmware upgrade for just about all iFi audio digital products.


WHAT IS THE iFi GTO™ DIGITAL FILTER?

The Gibbs Transient Optimised filter (GTO) is named after the ‘Gibbs phenomenon²’ in mathematics.

Wikipedia referred to the Gibbs phenomenon as “the peculiar manner in which the Fourier series of a piecewise continuously differentiable periodic function behaves at a jump discontinuity. The n th partial sum of the Fourier series has large oscillations near the jump, which might increase the maximum of the partial sum above that of the function itself. The overshoot does not die out as n increases, but approaches a finite limit.”

Most crucially, this is one cause of ringing artefacts’ in signal processing which the GTO addresses.

Way back in May 2011, the parent company of iFi audio, AMR, pioneered an earlier version of this filter in the DP-777 digital processor where it was available as an ‘Organic’ filter. Since 2011, more time has been invested into producing a filter that offered both better compatibility and technical performance than non-oversampling, while delivering a transient optimised performance that differs as little from non-oversampling as possible, delivering the new GTO™ filter.



Non-oversampling Transient response vs Organic Digital Filter AMR DP-777

Used by permission, Stereophile: <https://www.stereophile.com/content/ifi-audio-pro-idsd-da-processorheadphone-amplifier-measurements>

- 1) The original iDAC micro cannot receive this upgrade
- 2) https://en.wikipedia.org/wiki/Gibbs_phenomenon

Tech Note

No doubt there will be extended debate if our GTO™ digital filter offers the right trade-off, compared to others. To us the two key qualities we sought was to shape of the unavoidable transient or time domain distortion so that is free of any ‘pre-ringing’ and that completes its impulse response within a fraction of the Haas (precedence effect) window; to remain in effect, inaudible to the human ear.



What we really refer to when we are talking about ringing in digital filters is actually a form of ‘Echo’ or ‘Reverb’ where, in addition to the actual transient time-shifted lower amplitude, copies of the impulse are generated using delay lines (see also the *transients and digital filters* section later on).

The human hearing itself is subject to an inherent transient post (impulse) ringing that completely decays within around 0.7mS³ (see also the *transients and the human hearing* section later on).

The GTO filter’s transient post-ringing decays completely within 0.72mS for a 44.kHz source, ensuring that the unavoidable blurring of the transient response cannot be heard, but is integrated by the human hearing into the original transient.

This is in stark contrast to some alternative filter concepts. For example, the ‘Transient Aligned’ filter seeks a *maximum number of taps*, leading to an impulse response that falls well outside the Haas window. ie. its ‘ringing’ is very audible, in part because there is a pre-ringing (or pre-echo) present and in part through the sheer length of the delay line used.

For example, the 16k tap Transient Aligned filter in the Pro iDSD has an impulse response with equal pre- and post-ringing trail of around 186mS @ 44.1kHz sample rates, or a total 386mS worth of ringing. This is certainly sufficient time delay to be perceived as reverb. Using an even larger number of taps lengthens this impulse response even more.



Transient Aligned Digital Filter Transient response vs GTO™ Digital Filter iFi iDSD Pro

Used by permission, Stereophile: <https://www.stereophile.com/content/ifi-audio-pro-idsd-da-processorheadphone-amplifier-measurements>

It may be of course, that some will prefer the sound of a very long filter, with large amounts of ringing/reverb/echo as the result is often perceived as extra added spaciousness, however, to anyone seeking to be close to the original musical performance such additives

3) “Response of the human tympanic membrane to transient acoustic and mechanical stimuli: Preliminary results” Payam Razavi, Michael E. Ravicz et al - Hear Res. 2016 Oct; 340: 15–24.

are usually unwanted.

In the end, with the iFi GTO filter, by keeping the filter short and without pre-ringing, the filter response is inaudible because it is masked by the limits of the human hearing system. At the same time this filter still permits significant attenuation of unwanted ultrasonic images, compared to non-oversampling and also other attempts at “low tap number digital filter”.

Analogy: if a 20million mega pixel camera was used to take a picture of a straight line, the naked eye would see only a straight line. As the resolution is 'beyond' that of the human eye, any ultra-fine imperfections are not 'seen'. This is the same as with the GTO filter with human hearing.

If the GTO™ digital filter is so ‘perfect’, why include the other filters with the Pro iDSD? As remarked before, individual listeners may have different listening preferences and rather than imposing one option, even if we feel this option is not the best, we prefer to leave the choice down to the individual.

HOW TAPS RELATES TO WHAT IS HEARD

So far, we have identified that we prefer the GTO filter because it has few taps.

Because:

More taps = more reverberation.

Few taps = minimal reverberation


Reverberation⁴ is artificial. Sound engineers add reverb to make recordings more spacious, artificially so. Digital filters introduce reverb by the nature of their operation. In fact, a digital reverb unit operates precisely like a digital filter in principle – as depicted in this diagram.

Within digital filters are Digital Delay Lines which is defined by Wikipedia:

"A digital delay line is a discrete element in digital filter theory, which allows a signal to be delayed by a number of samples.

Delays of N samples is notated as $\{z\}^{-N}$ motivated by the role the z-transform plays in describing digital filter structures.

Digital delay lines are widely used building blocks in methods to simulate room acoustics, musical instruments and digital audio effects."



To our ears, the GTO filter simply sounds ‘right’ without any hint of artefacts or exceptional detail that feels ‘processed’, by avoiding large number of tap’s that add excessive reverb.

⁴ https://en.wikipedia.org/wiki/Reverberation#Digital_reverberators and https://en.wikipedia.org/wiki/Digital_delay_line




Tech Note

TRANSIENTS AND THE HUMAN HEARING

The human ear is a marvelous system with an incredible dynamic range (~135dB in middle frequencies) huge bandwidth (almost 1:1000) and a transient resolving ability that exceeds the upper limit of hearing steady state tones. Yet it is also subject to limiting factors that result in, so to speak, “blind spots” in its behavior that do not exist in purely mechanical systems (e.g. microphone). These “blind spots” can mask some behavior which objectively is distortion to be inaudible. For example, harmonic distortion masking has been well documented since the at least the 1950s if not earlier and it is reasonable to consider that ‘ringing’ on transients is also masked to a certain degree.


If we wish to produce audio gear that is capable of operating in a way that subjectively sounds undistorted to the human hearing (the most logical preference), we must understand its limitations and capabilities. Here, we focus on the time-domain capabilities.

It has been shown that the human hearing’s time domain resolution for the initial transient may be as small as 5 μ s. Some debate remains as to the exact limits, though work done by Dr Peter Lennox of Derby University suggests a median between 13...18 μ s, or a location accuracy of less than 2 degrees.



Additionally, the transient response of the human hearing includes 500...700 μ s ringing caused by the ear’s mechanical system (Tympanic Membrane, Malleus / Incus / Stapes).

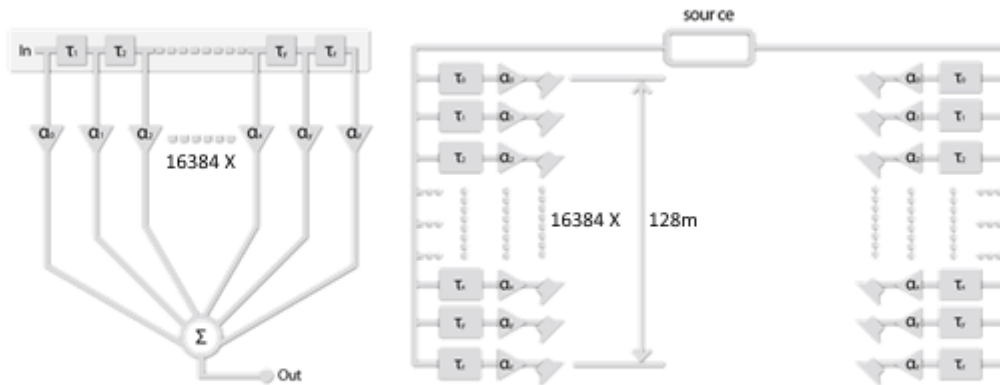
This ringing occurs after a transient event, there is no pre-ringing. The ringing in the ear’s system will mask any similar external ringing, which will instead be integrated into the transient, so it is inaudible.



Any pre-ringing is not masked by the human hearing, nor is any ringing that continues substantially beyond 500...700 μ s.


TRANSIENTS AND DIGITAL FILTERS

Digital filters introduce time-domain distortions that are unavoidable and a result of the inherent functional principle of digital filters. Digital filters are formed from delay-lines with so-called 'taps' usually at one audio sample intervals. Summing the delayed signals with varying gain, positive and negative, forms the filter function. Time-domain distortion, just as amplitude domain distortion, is accumulated along the recording and playback chain and cannot be removed unless its exact nature is known.




FIR type digital filter structure

Acoustic effect of FIR filter




Filter Response of FIR filter



Time domain of FIR filter

Used by permission, Stereophile: <https://www.stereophile.com/content/ifi-audio-pro-idsd-da-processorheadphone-amplifier-measurements>

Among others, the mathematical consequences of this filtering have been described by J. Willard Gibbs after whom the Gibbs phenomenon is named. A square-wave filtered perfectly with a filter that allows only the fundamental to pass, will be a sinewave. Transient distortion, however, will be very high. If a square-wave is filtered but not with a perfect filter that only allows the fundamental to pass, ringing is observed. No filtering restores the perfect wave.




square-

Tech Note

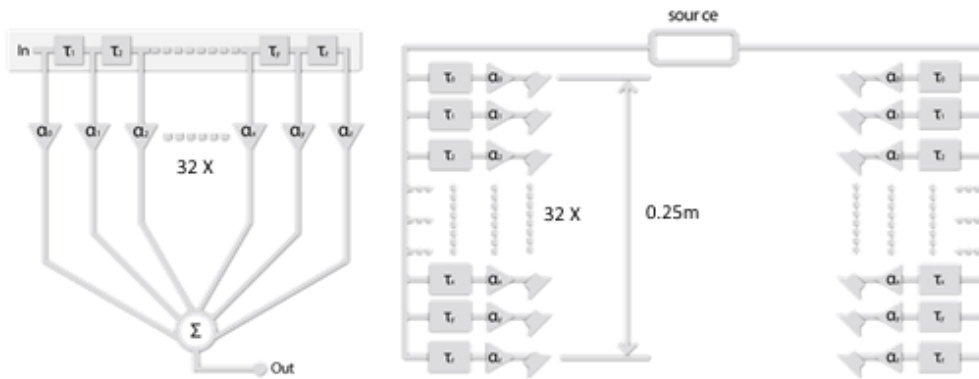
OPTIMISING DIGITAL FILTERS FOR TRANSIENTS

Considering the responses of the human ear, a digital filter can claim to offer a transient response distortion that does not materially alter the perception of music should have a transient response that has been completed in under 700µs and be free of pre-ringing.

Length of transient response of digital filters (shorter is better)



Such a filter could be an asymmetrical FIR filter with no more than 32 taps at 44.1/48kHz sample rate and no more than 64 taps at 88.2/96kHz sample rate, which is precisely what iFi have implemented as GTO Filter.

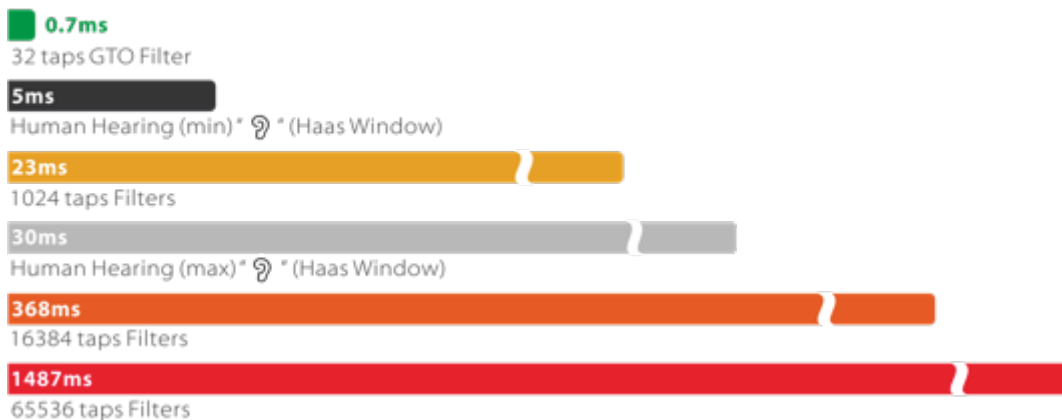


GTO type digital filter structure

Acoustic effect of GTO filter

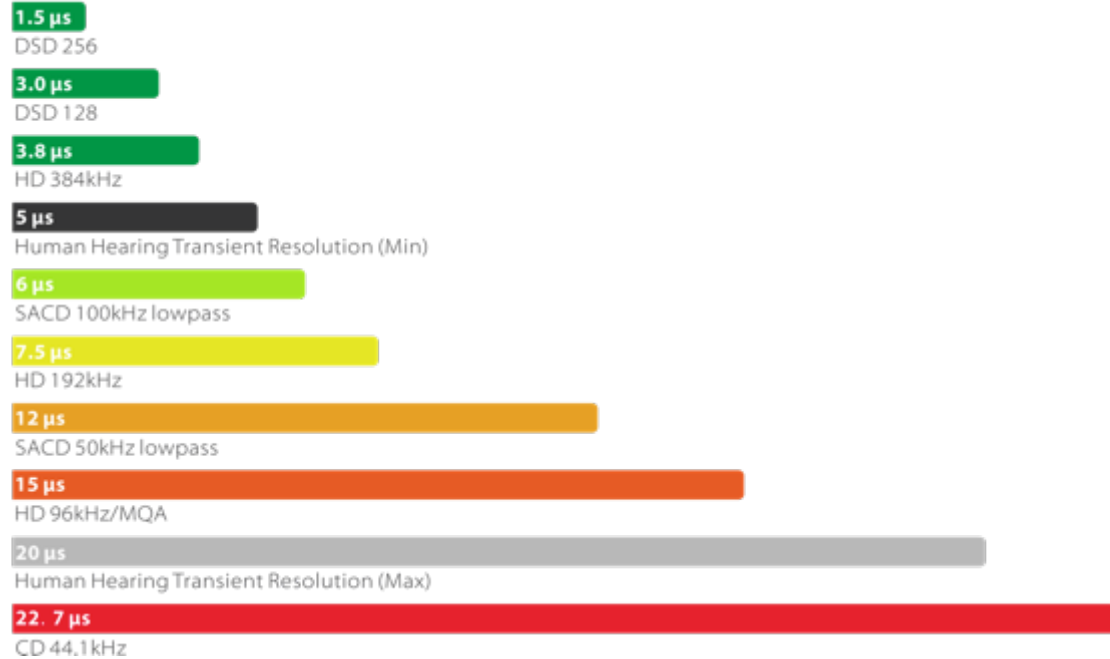
Given the 'Haas Window' within which reverb or ringing is integrated into the main arrival is agreed to be at minimum 5ms (some debate remains to the actual number), one may expect filters that require more than 5ms to complete their impulse response (or have in excess of around 220 Taps at 44.1kHz sample rate) to have the potential to be significantly audible as alteration.

Length of transient response of filters (shorter is better)



By contrast a 5µs transient resolving ability of the ear would require a sample rate of around 200kHz to ensure a transient response, while a 20µs transient resolving ability of the ear would require a sample rate of over 50kHz, with a good middle ground represented by a 88.2/96kHz sample rate.

Temporal resolution of formats (shorter is better)



The above suggests that at 44.1kHz sample rate digital audio lacks time domain resolution to be transparent and the common ‘long’ symmetrical FIR digital filter will have the potential to cause audible problems. It is thus desirable to use a higher sample rate and ‘short’ asymmetrical digital filters to ensure recording and playback are audibly transparent.

At this point in time even higher sample rate recordings tend to utilise “long” symmetrical FIR Filters, so their benefit is diluted. That said, if we know the filter response applied at the time of recording, the time-domain distortion in the recording could in theory be reversed and replaced by one that is not audible, this action is asserted for the complete end to end MQA coding system promoted by MQA Ltd, which is now supported by all current iFi products.

For recordings not available in MQA derived from analogue Masters or genuine high-resolution digital masters with a known provenance, selecting iFi’s GTO filter presents the highest possible fidelity.

HOW iFi HAS IMPLEMENTED THE GTO™ DIGITAL FILTER

A digital filter must be implemented somewhere, perhaps on an FPGA like the Xilinx Spartan 6 as used in the Pro iDSD, a DSP chip like an Analogue Devices Sharc or Blackfin. Recent developments on the XMOS platform have led to our becoming very familiar with offerings from the new XCore200 and the onboard DSP capabilities baked-in.



Tech Note

In the Pro iDSD we choose to use the XILINX Spartan 6 FPGA which also performs the DSD 1024 remastering, simply because the 16-Core XMOS is substantially loaded already.



Using any option requiring dedicated hardware would mean this filter can only be made available on products in development at higher price-points as FPGAs or DSP chips sufficiently capable are costly or on products using the xCore-200, which is on-board of the Pro iDSD.

GTO TRICKLE-DOWN

Unlike the xCore-200 series which include built-in DSP capabilities, implementing MQA on xCore-100 required the ground-up design of a DSP engine. With the DSP engine now available, our natural inclination was to have our new filter baked onto the XMOS Firmware.



Technical assistance came from an unexpected source; MQA. Using very different filter parameters as part of a more complex system including the decoder section requires serious DSP chops. The implementation of MQA rendering on the original xCore-100 series XMOS processors widely used in the iFi products required

extensive DSP work.

This took some extra work, and some fine-tuning of resource allocation, however with the version 5.3C of our XMOS firmware we now offer this new, Global Transient Optimised digital filter to all our customers, including just about all our legacy products, as an alternate option to the filters implemented in the DAC chip that apply for version 5.30 and before.

TECHNICAL BACKDROP

The 'Gibbs Transient Optimised' filter was developed by iFi according to our specifications in conjunction with the MQA team. We must make clear that GTO is not directly related to filter types used by MQA, it is not "MQA through the backdoor", but instead what we feel is the optimum solution for the playback of digital audio that has not undergone the MQA process. We would like to thank MQA for their technical assistance in integrating this into our firmware.