

OS/NOS Subjective Divide Investigation – Concluding Report

1. ACKNOWLEDGMENTS

I first wish to thank, Hans Polak. He was the primary technical interface with ZB, author of the PGGB software interpolation-filter, and managed preparation of all test file sets except for those of the 'Echo' experiment. In fact, it was Hans who suggested that we produce a common controlled set of test files for everyone to use. Thank you, for being tireless in your able assistance. The investigation certainly would not have been as successful or useful without your involvement.

I also wish to thank, Marcel van de Gevel for the FIR filter 'Echo' experiment. Which was conceived, produced and conducted by him. His analytical skills were utilized to produce a professional statistical analysis of the results. Outstanding work, as always, Marcel.

Appreciation, to our thread contributors, and experiment participants for patiently remaining with the thread all the while. You showed great perseverance in conducting the listening experiments, some of which were quite tedious. Without you, the investigation would have gone nowhere. Tipping of my hat, to all.

Finally, I wish to thank, Zaphod Beeblebrox, aka, ZB. Author of the PGGB software interpolation-filter. He quietly supported Hans and me from behind the scenes. Producing upsampled files simply at our request. ZB was always helpful, and never so much as hinted that we promote his worthy creation. For the record, the thread announcement of his PGGB-RT product was my idea. I did so, because our experiments had shown the technology to be worthy, and because it was being offered for free use.

2. OBJECTIVE

Determine why typical oversampled (OS) playback and non-oversampled (NOS) playback do not sound the same when it seemed that they should. This question is based on the fact that whether OS or NOS, digital audio image-bands are ultrasonic and not audible by human ears in either case.

An unplanned follow-on goal presented itself with the apparently successful satisfying of goal 'A'. Which was; to determine whether there was a means of CD music playback superior to either NOS or typical-OS based on the learnings of goal 'A'.

3. SUMMARY

3.1 Perhaps, the first thing to note about the investigation is that, logically, it was not applicable to those who do not hear a characteristic difference between typical-OS playback, and NOS playback. This subjective difference was the impetus for the investigation. It's also important to understand that the investigation was NOT founded on any assumption regarding whether OS or NOS was subjectively superior to the other. It was recognized that each exhibit characteristic subjective advantages. Rather, it was sought to identify why the two sound different when, on the surface, it seemed they should not. The notion that they should essentially sound the same stems from the fact that there are two fundamental objective differences between OS and NOS. By NOS, I mean, does

not utilize an OS FIR digital interpolation-filter. Whatever is the sample rate input to the DAC is the conversion rate out of it.

3.2 The objective differences between OS and NOS are well understood. One is the infamous NOS treble response droop. Which reaches -3.17dB @ 20kHz from a 44.1kHz sample rate, and is due to a gently curving frequency response masking effect that itself is caused by the stair-stepped Zeroth-Order-Hold (ZOH) output of nearly all audio DACs. It bears mentioning, that all DACs featuring ZOH converter operation, including OS DACs, exhibit this phenomena. The severity of which is a function of the output conversion rate, and the audio signal frequency. This is not of concern for OS-DACs with built-in digital filters because they convert samples at a high OS rate, which lessens the problem, and because they routinely feature digital EQ of the residual droop.

The second objective difference is how the two handle the image-bands, which are simply a fundamental consequence of the sampling process. Image-bands are copies of the desired signal-band which repeat ever higher in frequency. Image-bands are prevalent with both NOS-DACs and OS-DACs. Digital audio from 44.1kHz and higher, exhibit image-bands which are ALL ultrasonic and inaudible to the human ear. In other words, the human ear itself is acting as a low-pass filter which suppresses the image-bands to below audibility. The difference is that the image-bands are far better suppressed by OS interpolation-filters. Either way, the image-bands are suppressed to the level of being inaudible to middle-aged and older listeners. Especially, male listeners. Their inaudibility indicates that the desired audible signal-band is subjectively reconstructed, and should therefore should sound the same for NOS playback as it does for OS playback.

The investigation began in earnest with the assembling of a list of suspect causes. This list then formed a map to guide the investigation. The items on this list lent themselves to grouping in four categories. The subsequent process of suspect elimination was based on a combination of research, group listening experiments and logical deduction.

3.3 With the founding mystery established, we began by examining the item within category 'A'. The infamous NOS treble droop which was described above. Some NOS-DAC owners relayed that they don't feel that they perceive the droop. Which may have been due to age related HF hearing loss, but that's only speculation. We came to eliminate this item from suspicion, due to anecdotal listening reports that key OS/NOS subjective differences persist after analog EQ of the droop.

The next category to receive scrutiny was, 'B' Image-band handling. This would prove to be the most involved suspect category to investigate, as it had the greatest number of suspect items. Yet, all those items orbited around the question of whether common-place OS interpolation-filters introduce audible artifacts. A series of collective listening experiments was conducted to answer that question.

3.4 The first listening experiment had the objective of exposing whether a given OS interpolation-filter was audibly transparent. In other words, whether some OS interpolation-filters were introducing audible artifacts depending on how their processing was being performed. The experiment removed the possible confusion of artifacts due to the output sample rate simply being higher, with any due to the design and implementation of the interpolation filter.

The experiment was dubbed the 'Up/Down' experiment. It utilized four 44.1kHz source files, each source file was upsampled, and then down-sampled back to the original 44.1kHz rate. It was reasoned that an audibly transparent resampling process would result in a 44.1kHz resampled file

that sounded identical to the 44.1kHz source file. It was also reasoned that a NOS-DAC would be preferred to conduct the experiment because an OS-DAC's resident interpolation-filter which was less than transparent would cloud the results.

The second listening experiment was intended to reveal whether a very performance interpolation-filter would erase the subjective difference between OS and NOS. Participants were instructed to compare a set of 44.1kHz source files played NOS, to a set of those same source files upsampled to 88.2kHz via a high-performance interpolation-filter, the PGGB. Experiment results subsequently implicated the 'Image-Band Handling' category as containing the culprit cause for which we had been searching.

Experiment three, was an interesting attempt to identify the specific item within category 'B' which was the causal mechanism behind the difference phenomena. Known as the 'Echo' experiment, it was designed to determine whether listeners could detect the injection of signal echo/reflection exhibited by a certain type of ubiquitous on-chip interpolation-filter. Equiripple type FIR filters.

An unplanned fourth experiment was conducted for the purpose of determining whether 176.4kHz up/oversampling would produce playback which was subjectively superior to 88.2kHz up/oversampling.

3.5 Below, is the suspect list, including a brief explanation of the final disposition of the suspect items, shown in red:

A) FREQUENCY RESPONSE ERRORS:

=====

1) SINC aperture based -3dB droop @ 20KHz.

The Zeroth-Order-Hold (ZOH) based, -3dB droop @ 20KHz treble droop produced by NOS DACs converting 44.1kHz rate signals is easily corrected with analog equalization. Reports from those having built the Abraxalito NOS DAC, was that the NOS specific characteristics which were apparent before EQ, persist after EQ. Which also has been my anecdotal experience.

Therefore, the ZOH treble droop was judged as being unlikely responsible for the subjective difference between OS, and NOS. While anecdotal reports such as these do not constitute proof, the number of properly controlled experiments it would have taken to acquire proof was not practical.

B) RECONSTRUCTION/IMAGE-BAND HANDLING:

=====

2) Lack of an FIR interpolation-filter, freeing the DAC from certain processing 'artifacts' such as:

The investigation in to suspect category group 'B2' resulted in certain OS interpolation-filters as being the cause of NOS and OS, sounding characteristically different. Perhaps, not coincidentally, these are typically on-chip implementations. Of, on-chip OS filters, nearly ubiquitous are half-band, equiripple linear-phase implementations. Which was the exact causal mechanism from the items below was not determined for reason of practical experimental limitation.

a) time-domain signal echoes/reflections produced within equiripple on-chip FIR filters.

A perceptual test was conceived and executed by Marcel van de Gevel, in an attempt to identify whether equiripple filter type signal echo was the root cause mechanism. Known as the 'Echo' experiment, the results were deemed inconclusive, partly because test respondents overwhelmingly utilized OS DACs to conduct the experiment. That resulted in all test files being contaminated by the very type of echo error phenomena which was being tested for detection.

b) impulse-response ringing

While not specifically tested, this particular suspect can be logically be removed from guilt, because the PGGB 88.2kHz upsampling experiment indicates that a high-performance OS filter essentially sounds the same as NOS. The higher performing is an interpolation-filter, the longer is the filter's impulse-response ringing. Therefore, logically, impulse-response ringing is not the cause of the OS/NOS sound difference.

c) half-band filters violating Nyquist

Half-band filters are always down only -6.02dB at the Nyquist frequency. Which has serious implications for such filters utilized in ADC service, where in-band aliasing could result. This becomes a different level of concern for DAC image-band suppression, where the lowest frequency affected by the filter is still ultrasonic. Even without any electronic suppression, image-frequencies don't extend below the DAC output conversion rate's Nyquist frequency.

d) clipping on peak sample normalized recordings - the intersample-overshoot issue.

Driven by the, 'Loudness War'. This is an issue with recordings that reach, or closely reach digital full scale. Such digital signals can drive non-pre-attenuated interpolation-filters to clip. The question was, whether this was responsible for the OS/NOS subjective difference. Logically, the answer is, very probably not, as the characteristic difference between OS and NOS is apparent via most recordings.

e) analog image-band suppression inherently sounds different than digital suppression?

Since high-performance OS image-band suppression essentially removes the subjective difference between NOS and OS, we logically conclude that purely analog image-band suppression is not necessary for doing so. Therefore, whether or not analog image-band filtering sounds different from digital filtering, isn't the reason why typical-OS filters sound different from NOS.

3) Imperfect signal reconstruction due to insufficiently suppressed image-bands. In other words, because the signal waveform is may not be sufficiently reconstructed per the sampling theorem.

Whether the theoretically more accurate signal reconstruction produced by longer impulse-response ringing is perceptible from within the desired signal-band, remains something of a mystery. In any case, since the long filter of high-performance OS sounds very similar to having no filter at all of NOS, we logically deduced that reconstruction accuracy was not responsible for the subjective difference.

4) The unsuppressed image-bands producing audible IM products directly within the ear.

While research on this item appeared very sketchy, IMD seems highly unlikely to produce the relaxed character associated with NOS. In fact, just the opposite reaction would be expected. In addition, anecdotal reports were that high-order analog filtered image-band DACs retained the classic NOS character even while suppressing any potential IMD generating image-frequencies.

C) ALTERED JITTER IMPACT

=====

5) Different jitter impact due to fewer D/A conversion cycles per second.

6) Reduced supply and ground noise due to slower clock rates.

While these are undoubtedly factors affecting the sound quality of a DAC, the results of the 88.2KHz upsampling experiment indicated that typical-OS interpolation-filters are either primarily, or completely responsible for the subjective difference in OS/NOS character. Therefore, all of category 'C' was removed from significant suspicion.

D) SAMPLE-PERIOD RELATED QUANTIZATION ERRORS

=====

7) Converter settling-time becomes a smaller percentage of each conversion period as the conversion rate is made slower.

8) Harmonic-distortion may be sample-rate dependent.

Category 'D' was removed from significant suspicion for the same reason as category 'C', the results of the 88.2KHz upsampling experiment. Which is not to suggest that there isn't opportunity for worthwhile improvement in this category, for both NOS-DACs and OS-DACs. Such as, DAC and I/V circuit settling-time error mitigation. The two primary means of which are to, passively RC filter it or utilize an active Sample-and-Hold circuit to gate it.

4. LISTENING EXPERIMENTS

4.1_44.1kHz up/down resampling filter transparency:

Hans' up/down resampling experiment was intended to reveal whether a given interpolation-filter performed audibly transparent rate conversion, with other sample rate dependent phenomena were kept unchanged.

4.2 88.2kHz upsampling filter preference/equivalence:

Participants listened for a subjective preference between a set of 44.1kHz source files played NOS, and those same files upsampled to 88.2kHz by the PGGB high-performance PC/Mac based interpolation filter, before being sent in to the NOS-DAC.

4.3 176.4kHz upsampling filter preference/equivalence:

Participants listened for a subjective preference between a set of 88.2kHz and 176.4kHz files, both sets upsampled from the same 44.1kHz source files by the PGGB interpolation-filter.

Up/oversampling Experiments Merged Results Document. Post #1773:
<https://www.diyaudio.com/forums/digital-line-level/371931-makes-nos-sound-178.html#post6775491>

4.4 The FIR filter signal 'Echo' detection experiment:

The Echo experiment was designed to test a hypothesis about a leading root cause candidate underlying the NOS/OS subjective divide. Identified by R. Lagadec and T.G. Stockham all the way back in 1984, at the first years of CD. They discovered that a certain type of FIR interpolation-filter designed utilizing the McClellan/Remez exchange algorithm method, known as an equiripple filter. Equiripple filters suffer parasitic signal reflection/echo in the time-domain. These echoes are not to be confused with the familiar FIR filter SINC-function based impulse-response ringing.

Such echoes are objectively evidenced by the very low level and repeating frequency-domain ripples which span an equiripple filter's passband response. The echoes are also low in level, and not obvious like the echo from shouting in a hall. These tiny passband ripples are commonly visible in equiripple type interpolation-filter data sheet performance graphs. The great majority of DAC chips which feature on-chip OS interpolation-filters continue to be equiripple designs. The Echo experiment was conceived, conducted and results analyzed by, Marcel van de Gevel. He utilized the 'Gold-Wave' sound editing tool to inject simulated echo phenomena in to controlled sets of test files. It was postulated that the injected echo would be subjectively detected by listeners.

'Echo' Experiment Results Document. Post #1518: <https://www.diyaudio.com/forums/digital-line-level/371931-makes-nos-sound-152.html#post6761338>

5. 'FINDINGS' DISCLAIMER

From the beginning, it was well understood that we were not running an AES approved research experiment. As hobbyists, we were forced within hard practical limitations. One works with what one has. Otherwise, I suppose, we should all just shut up, and accept whenever someone tells us that a given component is perfect, simply because whatever specifications are available dictate that it MUST be perfect.

I've intentionally used words such as; "appears", or "seems", or "indicates", when referring to the supposed findings of our listening experiments. This is a very intentional acknowledgement that they are not scientific, being largely uncontrolled and feature unique listening circumstances and playback systems for each participant. This was a practical reality, of being a voluntary group of internationally dispersed hobbyists, having limited (non-commercial) resources and conducting in-home listening experiments.

In addition, the very low tally of experiment responses prevents any statistical confidence in the findings. Who know, more test responses may have built an alternate finding, or instead, may have reinforced the apparent findings. While these factors necessarily render the finding questionable, that is not the same as automatically rendering them incorrect. Since it would be too cumbersome to constantly restate this disclaimer throughout the report, suffice to say, that the inherent scientific faults of the experiments are fully acknowledged and not disputed.

6. WHAT WAS FOUND

Common-place (typically, on-chip equiripple type) OS interpolation-filters introduce audible artifacts which are absent from either NOS or from high-performance (typically, PC/Mac resident, software) OS.

The PGGB 88.kHz upsampling experiment indicated that OS is not inherently the cause of the subjective difference from NOS, but rather, somehow insufficiently performed OS. The question naturally arises, is NOS actually the way digital playback should sound after all? While we didn't seek to answer that intriguing question, we did find that PGGB 88.2kHz upsampling sounded a lot more like NOS, and a lot less like typical-OS. With both high-performance-OS and NOS sounding more like each other, typical-OS is implicated as the one which sounds the more incorrect.

Resulting DAC topologies are to either; Utilize high-performance OS, which doesn't introduce audible artifacts. Or, don't implement any OS at all, and simply go pure NOS - optionally, with analog treble EQ. A third topology is suggested by the findings, however, since it was not evaluated, it is not included among the suggested options in this report.

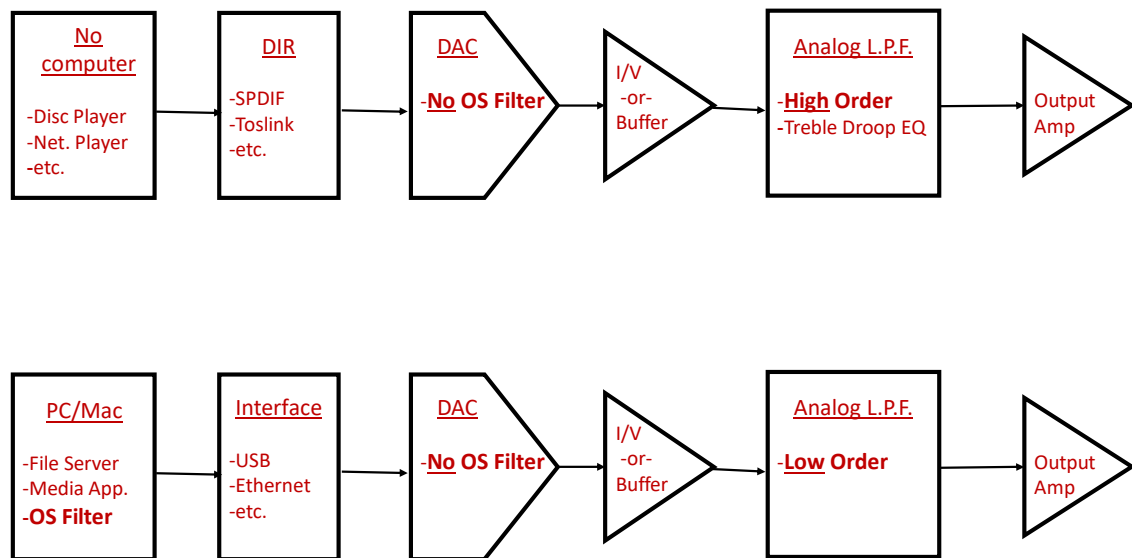
7. WHAT WAS NOT FOUND

We did not identify the root cause mechanism which introduces audible artifacts via typical OS interpolation-filters. The potential mechanisms are several. Concern over respondent fatigue with the quantity of permitted one attempt to test a particular mechanism. Which was the Echo experiment.

8. SUGGESTED APPLICATION OF FINDINGS

Two basic DAC topologies emerged as a result of the investigation. One for non-computer sourced playback, and one for computer sourced playback. Depicted on the next page.

The potential third topology mentioned above is based around some programmable chip solution. DSP or MCU or FPGA. Which would re-take the topological position which dedicated OS-filter chips used to occupy before their function was integrated on to the DAC chip.



8.1 **Non-computer sourced** music systems, there would be no FIR interpolation-filtering applied anywhere. If utilizing a DAC IC featuring an on-chip interpolation-filter, as do most all contemporary devices, the digital filter unit must be capable of being passed. In other words, this is a classic NOS-DAC.

Simulations showed that a high-order low-pass filter can attenuate a 25kHz image-frequency by 16dB while maintaining flat response to 20kHz, including EQ of the 44.1kHz NOS treble droop. Anecdotal reports of the subjective effect of high-order analog image-band suppression filters were that they remove a certain 'dirty' quality produced by some instruments, such as the upper range of piano. Otherwise, they leave the NOS character fully intact. In terms of objective performance, analog filters fall well short of digital filters, and also are sensitive to passive parts tolerance. Where as digital filters don't utilize passive parts.

8.2 **Computer sourced** music systems, in addition to music file server duties and music content management, the computer would perform up/oversampling via software based high-performance interpolation-filter. There would be no be other interpolation-filter within the playback chain. Which means, the core DAC, again, is operating in NOS mode. Even so, this still constitutes OS playback. Only, it is performed external to the DAC.

The high-performance interpolation-filter utilized in our experiments was the PGGB. Available at: www.remastero.com. One of the major advantages of locating the interpolation-filter in a computer is that, OS filter technology improvements are readily added to the existing system as they become available. The highest performing software filter setting within Audacity, which I believe is based on the 'SoX' utility, was also evaluated but judged inferior to the PGGB in our first collective listening experiment. Additionally, we received anecdotal reports of very good subjective results via the, 'HQPlayer'. Which is another software interpolation-filter. Due to certain practical concerns, it was

decided not to have thread participants perform an experiment to determine whether the HQPlayer was subjectively superior to the PGGB. So, that remains an open question.

Inclusion of a post DAC, low-order analog filter should be considered. This would only be for the purpose of protecting downstream amplification from OS image-band spectra, which might otherwise provoke intermodulation distortion. Simulation of a passive 3rd order, low-pass LCL filter predicted 23dB attenuation of the first (lowest) image-band frequency produced by 176.4kHz upsampling rate, which is 154kHz. In addition, 176.4kHz rate ZOH treble droop lessens to -0.2dB @ 20kHz. As a result, analog EQ of this small droop is likely not necessary, but still possible by slight peaking the response of the 3rd order filter if desired.

9. AREAS FOR FURTHER INVESTIGATION

9.1 Determine the root cause mechanism behind why typical on-chip interpolation-filters introduce subjective artifacts.

9.2 Determine whether DAC settling-time errors are audible, and if so, determine the efficacy of different mitigation solutions. Such as RC filter versus Sample & Hold circuit.

9.3 Determine why the Belafonte 'Day-O' source file in the first experiment, the only source file which was directly digitized from vinyl LP (by, Hans Polak) was unanimously preferred over the PGGB up/down resampled alternative file. This was the only 44.1kHz source file among four, to be unanimously chosen over the resampled 44.1kHz respective alternative.

10. META LEARNINGS

10.1 Perhaps, the most significant learning was how difficult it would be to conduct scientifically controlled, statistically confident listening experiments utilizing only a volunteer collective of remotely located hobbyists. This difficulty necessarily renders such experiment's findings as always questionable. Certain limited controls can be realized, which in our case, was the production and distribution of identical, controlled sets of test files. Distributed via Dropbox download.

In addition, the fact that participants were hobbyists meant that experiments needed to be limited in inconvenience, and in quantity so as to persuade voluntary participation. One particularly perplexing fact was that at a time when the thread was approaching 80K views, we were having difficulty obtaining six experiment responses. I speculate that experiment fatigue was beginning to become a factor. Perhaps, it was a combination of factors. Complicating the conducting of the experiments, was that the investigation organically and rapidly blossomed to become much more than was originally envisioned. As a result, clear and coherent instructions for an experiment had, sometimes, not always been thoroughly thought through at the time the experiment was introduced.

10.2 All of which lead to the question of, whether hobbyist run collective listening experiments are useful at all? I think that the answer to that question depends on for what purpose the resulting 'findings' are to be used. They should not be used for dictating what attributes/features correctly constitute a given type of component. For the purpose of encouraging the consideration and exploration of non-conventional features or implementations, however, I feel that hobbyist

experiments do offer worthwhile value. For example, I unexpectedly discovered what a software interpolation-filter was subjectively capable of doing.

This report concludes my leadership of the thread. It will remain active, of course, but my purpose in launching it has been fulfilled to my satisfaction. Perhaps, someone will lead a pursuit in to one of the suggested areas for further investigation. In any event, it has been a memorable and fun experience.

So long, and thanks for all the fish. (Nod to, ZB)