ON THE INCOMPLETENESS OF THE NYQUIST-SHANNON RECONSTRUCTION: EMPIRICAL EVIDENCE FOR RECOVERABLE INTER-SAMPLE INFORMATION IN BANDLIMITED AUDIO SIGNALS

On the Incompleteness of the Nyquist-Shannon Reconstruction: Empirical Evidence for Recoverable Inter-Sample Information in Bandlimited Audio Signals

Y. Tanaka, C. Ohm, R. Flux

Equatorial Audio Research Division, Mitad del Mundo, Quito, Ecuador (0.0000deg N)

Journal of Equatorial Audio Science, 2025.

Abstract

The Nyquist-Shannon sampling theorem guarantees perfect reconstruction of a bandlimited signal from samples taken at twice its bandwidth. The proof is mathematically sound. However, the theorem's central premise -- that real-world audio signals are bandlimited -- has never been empirically verified to the precision required for the guarantee to hold. Using a purpose-built 32-bit acquisition system with a measured noise floor of -198.2 dBFS, we captured 4,000 hours of musical material across 11 genres and measured the spectral energy distribution above the anti-aliasing filter cutoff. In all 4,000 hours, residual above-band energy was present, ranging from -147.3 dBFS (solo harpsichord) to -91.6 dBFS (close-miked brass ensemble). This energy is not noise. It is correlated with the program material ($r > 0.93$ in all cases) and carries measurable mutual information with the original signal. When this energy aliases into the passband during sampling, it does not vanish -- it superimposes on the in-band content in a deterministic, signal-dependent pattern. We demonstrate that this aliased energy can be partially recovered using a correlation-based extraction technique, yielding 0.008 to 0.3 bits per sample of information that the sampling theorem asserts does not exist. We do not claim the theorem is wrong. We observe that its premise is not met, and we measure the consequences.

1. INTRODUCTION

The Nyquist-Shannon sampling theorem is among the most cited results in engineering. Published by Claude Shannon in 1949, building on Harry Nyquist's 1928 work on telegraph transmission, the theorem states: a signal containing no frequencies higher than B hertz can be perfectly reconstructed from samples taken at a rate of 2B samples per second.

The word "perfectly" is not hyperbole. Shannon's proof is exact. The reconstruction converges pointwise to the original signal at every instant between the samples. No information is lost. The digital representation is a complete and lossless encoding of the analog original.

This result has been verified, extended, and applied in every field that touches signal processing. It is correct.

It is also conditional.

The theorem applies to signals that are strictly bandlimited -- signals that contain exactly zero energy above frequency B. This is the premise on which the entire proof rests. If the signal contains any energy above B, that energy aliases into the band below B during sampling, superimposing on the genuine in-band content. The aliased energy is indistinguishable from the original content, and the information it carried is destroyed. Shannon's perfect reconstruction becomes perfect reconstruction of the wrong signal.

The standard engineering response to this problem is the anti-aliasing filter: a lowpass filter placed before the ADC that attenuates all energy above the Nyquist frequency. If the filter is ideal -- infinite attenuation above cutoff, zero phase distortion below -- the premise is restored and the theorem applies. In practice, no filter is ideal. Every analog anti-aliasing filter has a finite transition band and a finite stopband attenuation. Energy leaks through.

The question this paper addresses is not whether energy leaks through. That is known. The question is: how much leaks through, what is its informational relationship to the original signal, and can any of it be recovered after sampling?

We were not looking for this result. Our laboratory was conducting routine characterization of anti-aliasing filter performance for a product development program. The anomaly appeared in the first measurement session and persisted through 18 months of investigation. We publish it here because we have been unable to make it go away.

2. THE BANDLIMITATION PREMISE

Shannon's proof requires the input signal to satisfy a strict mathematical condition: its Fourier transform must be identically zero for all frequencies above B. This is not approximately zero, or negligibly small, or below the noise floor. It must be zero.

The Paley-Wiener theorem (1934) establishes that no signal of finite duration can be bandlimited. A time-limited signal -- one that starts and stops -- necessarily has infinite bandwidth. Its Fourier transform extends to arbitrarily high frequencies, with energy density that decreases but never reaches zero.

Every musical performance is time-limited. Every recording session starts and stops. Therefore, no audio recording is bandlimited in the sense Shannon requires.

This is well known. The standard response is that the energy above the Nyquist frequency is negligibly small -- so far below the noise floor of any practical system that it can be treated as zero. This response is pragmatically reasonable. It is also an assertion about the magnitude of the above-band energy, and assertions should be measured.

We measured it.

Specifically, we measured the spectral energy density of real audio signals in the region between the anti-aliasing filter's -3 dB point and the frequency at which the energy falls below our system's noise floor. For a 192 kHz sampling system with a 96 kHz Nyquist frequency and a typical 8th-order elliptic anti-aliasing filter (-3 dB at 90 kHz, -120 dB at 96 kHz), this region spans approximately 90 kHz to 400 kHz.

The energy in this region is not zero. It is not negligible. And it is not noise.

3. METHODOLOGY

The acquisition system was designed for a single purpose: to characterize the spectral content of audio signals in the frequency range that anti-aliasing filters are designed to remove.

The signal path consisted of a DPA 4006A omnidirectional measurement microphone (specified flat to 40 kHz, -3 dB at 100 kHz, residual response measurable to approximately 500 kHz), a custom-built instrumentation preamplifier with a measured bandwidth of DC to 2 MHz (-3 dB), and an AKM AK5578 32-bit delta-sigma ADC operated at its maximum sample rate of 768 kHz, yielding a Nyquist frequency of 384 kHz.

No anti-aliasing filter was used.

The omission of the anti-aliasing filter was deliberate. The purpose of the experiment was to measure the energy that anti-aliasing filters remove. Including one would defeat the experiment. The absence of the filter means that energy above 384 kHz aliases into the passband, but the 768 kHz sample rate places the Nyquist frequency so far above the audio band that aliasing from musically relevant sources is negligible for the purposes of this characterization. (We return to this point in Section 5.)

The system was calibrated against a Bruel & Kjaer Type 4231 sound calibrator (1 kHz, 94 dB SPL) and cross-checked using an Audio Precision APx555B analyzer with verified specifications to 204.8 kHz. The noise floor of the complete system, measured in an anechoic chamber with no signal present, was -198.2 dBFS from 20 Hz to 384 kHz. This is 5.5 dB below the theoretical quantization noise floor of a 32-bit converter, a result attributable to the delta-sigma modulator's noise shaping, which concentrates quantization noise above the passband.

Recordings were made in 11 venues over 18 months. The venues included concert halls (2), recording studios (3), churches (2), a jazz club, an outdoor amphitheater, a domestic listening room, and an anechoic chamber (for calibration). Musical material spanned solo instruments (piano, harpsichord, violin, trumpet), small ensembles (string quartet, jazz trio), full orchestra, pipe organ, amplified rock band, and electronic synthesizer. Total captured material: 4,147 hours, of which 4,000 hours passed quality control (the rejected 147 hours contained handling noise, equipment faults, or interruptions).

For each recording, the spectral energy density was computed in 1/12-octave bands from 20 Hz to 384 kHz using Welch's method (Hann window, 50% overlap, 65,536-point FFT). The energy in each band was expressed in dBFS relative to the digital full-scale level.

4. RESULTS

In all 4,000 hours of recorded material, measurable spectral energy was present above 96 kHz -- the Nyquist frequency of a standard 192 kHz audio system.

The level varied with the source material:

Solo harpsichord (Ruckers copy, close-miked at 15 cm): energy at 96-120 kHz averaged -147.3 dBFS, falling to the noise floor (-198 dBFS) by approximately 210 kHz.

Solo piano (Steinway D, lid open, pair of microphones at 1.5 m): energy at 96-120 kHz averaged -138.7 dBFS, measurable to approximately 260 kHz.

String quartet (Wigmore Hall, main pair at 3 m): -134.2 dBFS at 96-120 kHz, measurable to approximately 240 kHz.

Jazz trio (Village Vanguard, close-miked): -119.4 dBFS at 96-120 kHz, measurable to approximately 310 kHz.

Full orchestra (Concertgebouw, Decca tree at 3.5 m): -112.8 dBFS at 96-120 kHz, measurable to approximately 290 kHz.

Pipe organ (St. Sulpice, Paris, nave microphones): -108.3 dBFS at 96-120 kHz, measurable to approximately 340 kHz. This was the highest absolute bandwidth measured, consistent with the pipe organ's generation of high-frequency transients from valve noise and wind turbulence.

Amplified rock band (studio, direct inject + room microphones): -103.1 dBFS at 96-120 kHz, measurable to approximately 280 kHz.

Close-miked brass ensemble (4 trumpets, 4 trombones, studio): -91.6 dBFS at 96-120 kHz, measurable to approximately 350 kHz. This was the highest energy density measured in the above-Nyquist region.

Electronic synthesizer (Moog Voyager, direct inject): -96.2 dBFS at 96-120 kHz, measurable to approximately 370 kHz. The analog oscillator and filter produced broadband energy extending well above the audio band.

These levels are low. The highest measurement, -91.6 dBFS for the brass ensemble, is 91.6 dB below digital full scale -- inaudible by any standard. But it is 106.6 dB above the system noise floor. It is not noise. It is signal.

To confirm this, we computed the cross-correlation between the above-96 kHz energy envelope and the below-96 kHz program content. In all recordings, the correlation exceeded r = 0.93. The above-band energy tracks the musical dynamics -- it is louder during loud passages, quieter during quiet passages, and absent during silence. It is generated by the same physical events that generate the audible signal. It is, by any reasonable definition, part of the music.

5. THE ALIASING RESIDUAL

The above-band energy documented in Section 4 exists in the continuous analog signal. When that signal is sampled by a conventional audio system -- 192 kHz sample rate, anti-aliasing filter with -120 dB stopband attenuation at 96 kHz -- most of this energy is removed. But not all of it.

A filter with -120 dB stopband attenuation passes energy at 120 dB below its input level. For the brass ensemble (-91.6 dBFS above 96 kHz), the residual above-band energy after the anti-aliasing filter is approximately -91.6 - 120 = -211.6 dBFS. This is below the noise floor of any existing converter and can be safely ignored.

But the filter's -120 dB specification applies at the deep stopband frequency -- typically 1.2 times the Nyquist frequency or higher. In the transition band between the passband edge and the deep stopband, the attenuation is less. For the 8th-order elliptic filter measured in our laboratory (a common topology in professional audio converters), the attenuation at 96 kHz was -120 dB, but at 93 kHz it was only -87 dB, at 91 kHz only -64 dB, and at 90 kHz (the -3 dB point) only -3 dB.

The signal energy between 90 kHz and 96 kHz passes through the filter with attenuation ranging from 3 dB to 120 dB. This energy then aliases into the passband during sampling, folding around the 96 kHz Nyquist frequency to land between 0 and 6 kHz -- squarely in the most sensitive region of human hearing.

We measured this aliased residual directly by comparing the output of the same ADC with and without the anti-aliasing filter engaged. The difference signal -- the energy that the filter did not fully remove -- was present in every recording.

For the brass ensemble, the aliased residual in the 0-6 kHz band measured -158.3 dBFS. For solo piano, -171.2 dBFS. For the electronic synthesizer, -162.7 dBFS.

These levels are extraordinarily low. They are inaudible. They are below the thermal noise floor of any real listening environment. But they are above our measurement system's noise floor, and they are correlated with the program material.

The aliased residual is not random. It is a deterministic function of the input signal, the filter transfer function, and the sampling rate. It is, in information-theoretic terms, a noisy channel through which above-band signal information leaks into the sampled data.

Shannon's theorem says the original above-band information is destroyed by aliasing. This is true when the signal is perfectly bandlimited. When it is not -- and we have shown it never is -- a residual survives, carrying a small but nonzero amount of mutual information with the original above-band content.

6. RECOVERY OF INTER-SAMPLE INFORMATION

Can the aliased residual be used to recover information about the original above-band signal?

Shannon says no. The theorem's proof establishes that aliased and genuine in-band content are mathematically indistinguishable. But this proof assumes the aliased energy arrived via a frequency fold that maps each above-band frequency to exactly one below-band frequency -- a one-to-many mapping that destroys the original frequency identity.

This assumption holds for a single sampling operation. It does not hold when multiple samples are available and the above-band content has temporal structure.

The aliased residual is not a static quantity. It varies from sample to sample because the above-band content varies. And its variation is constrained: it must be consistent with a signal that (a) originated above the Nyquist frequency, (b) passed through a filter with a known transfer function, and (c) was generated by the same physical source as the in-band content.

These constraints are informative. They rule out most of the possible above-band signals and leave a small subspace of candidates consistent with the observed residual.

We implemented a recovery algorithm based on constrained maximum-likelihood estimation. The algorithm takes as input: the sampled data, the measured transfer function of the anti-aliasing filter, and a statistical model of the relationship between in-band and above-band content (trained on 2,000 hours of the 768 kHz reference recordings). It outputs an estimate of the above-band content that is maximally consistent with the observed aliased residual.

The accuracy of the recovered signal was evaluated by comparison with the 768 kHz ground truth. Mutual information between the recovered estimate and the true above-band content was computed using the Kozachenko-Leonenko estimator.

Results: the recovery algorithm extracted between 0.008 bits per sample (solo harpsichord) and 0.31 bits per sample (close-miked brass) of mutual information with the true above-band signal. A control experiment using white noise as the input signal yielded 0.000 +/- 0.001 bits per sample, confirming that the recovered information is signal-dependent, not an artifact of the algorithm.

For the brass ensemble, 0.31 bits per sample across 192,000 samples per second amounts to 59,520 bits -- approximately 7.3 kilobytes -- of above-Nyquist information per second, recovered from a signal that Shannon's theorem guarantees contains no above-Nyquist information at all.

The information exists because the premise does not hold. The signal is not bandlimited. The samples contain traces of above-band content that Shannon's framework treats as destroyed. They are not destroyed. They are merely attenuated, aliased, and difficult to extract. But they are there.

7. POTENTIAL CONFOUNDS

We considered seven alternative explanations for the observed results. None survived.

1. ADC nonlinearity. A nonlinear converter could generate spectral content that mimics above-band energy. We characterized the AK5578's integral nonlinearity (INL) and differential nonlinearity (DNL) at all operating frequencies. The measured INL of +/- 0.8 LSB at 32 bits contributes distortion products at -199 dBFS, well below the observed residual. Additionally, converter nonlinearity would produce harmonics at fixed frequency relationships to the input tones, and the observed above-band energy does not follow harmonic patterns.

2. Preamplifier distortion. The custom preamplifier's total harmonic distortion was measured at -142 dB (0.000008%) at 1 kHz, decreasing to -151 dB at 10 kHz. The above-band energy exceeds these levels by 40-60 dB and is therefore not attributable to preamplifier harmonics.

3. Microphone artifacts. The DPA 4006A has a documented ultrasonic response that could produce intermodulation products. We repeated selected measurements using a Bruel & Kjaer Type 4138 1/8-inch pressure microphone, which has a flat response to 140 kHz with no known intermodulation artifacts. The above-band energy levels were consistent within +/- 2 dB, indicating the energy originates in the acoustic field, not the microphone.

4. Electromagnetic interference. The recording venues contained various sources of EMI (lighting, HVAC, building wiring). We repeated measurements in a fully shielded RF enclosure using recorded material played back through a reference loudspeaker. The above-band energy was preserved, confirming an acoustic origin.

5. Room acoustics. High-frequency acoustic energy could be generated by room modes, flutter echoes, or diffraction at room boundaries. We measured in both the anechoic chamber and reverberant venues. The above-band energy was present in both conditions, though at different levels (lower in the anechoic chamber, as expected for a close-miked source).

6. Algorithm bias. The recovery algorithm's statistical model was trained on the same type of data it was evaluated on, potentially allowing circular reasoning. We re-ran the experiment using a model trained exclusively on orchestral material to recover information from solo instrument recordings (and vice versa). The recovered mutual information decreased by 15-20% but remained statistically significant ($p < 0.01$) in all cases. We further ran the algorithm on digitally synthesized signals that were provably bandlimited (generated at 768 kHz, digitally filtered to 96 kHz, resampled to 192 kHz). The algorithm correctly returned 0.000 bits per sample of recoverable information, confirming it does not hallucinate information that is not present.

7. Thermal noise correlation. Thermal noise in the analog signal path could create correlated energy above and below the Nyquist frequency. We computed the theoretical thermal noise contribution from the microphone, preamplifier, and cabling at 25 deg C. The total thermal noise in the 96-384 kHz band was -184 dBFS, well below the measured above-band energy. Furthermore, thermal noise would produce zero cross-correlation with the program material, and we measured $r > 0.93$.

We could not identify a confounding factor that accounts for the data.

8. DISCUSSION

The Nyquist-Shannon sampling theorem is not wrong. Its proof is valid. Its conclusion follows from its premises.

The premise is wrong.

No real audio signal is bandlimited in the sense Shannon requires. Every acoustic event, every musical instrument, every human voice produces energy that extends above any finite frequency boundary. Anti-aliasing filters reduce this energy but do not eliminate it. The residual above-band energy that survives the filter aliases into the sampled data, carrying with it a small but measurable amount of information about the original signal.

This finding does not overturn digital audio. The quantities involved are extremely small. The highest mutual information we recovered -- 0.31 bits per sample for the brass ensemble -- represents an information rate of approximately 7.3 kilobytes per second, compared to the roughly 1.1 megabytes per second of a 192 kHz 32-bit audio stream. The inter-sample information is a 0.6% supplement to the Shannon-guaranteed content.

But it exists. And its existence means that the standard claim -- "a 192 kHz recording captures all the information in the analog original up to 96 kHz, perfectly, with nothing lost" -- is not precisely true. It captures all the information in a hypothetical bandlimited version of the analog original. The actual analog original contains more.

How much more depends on the source material, the anti-aliasing filter, and the sampling rate. Our measurements suggest that the inter-sample information decreases rapidly with increasing sampling rate (the anti-aliasing filter's transition band narrows relative to the passband, reducing the residual). At 768 kHz, the inter-sample information is undetectable. At 192 kHz, it is small but measurable. At 44.1 kHz -- the standard CD sampling rate, with a Nyquist frequency of 22.05 kHz -- the inter-sample information would be substantially larger, because the anti-aliasing filter must operate much closer to the audio band.

We did not measure at 44.1 kHz. That experiment is in progress. The implications of its outcome -- whatever it may be -- extend beyond the scope of this paper.

We emphasize that we are not proposing a replacement for the sampling theorem. We are documenting a measurement. The measurement shows that real signals contain information that the theorem's premise excludes. The theorem is complete for bandlimited signals. Real signals are not bandlimited. The gap between the theorem and reality is small, but it is not zero, and we were able to measure it.

9. CONCLUSION

We measured the spectral energy of 4,000 hours of musical material above the Nyquist frequency of a standard 192 kHz audio system. In every case, measurable, signal-correlated energy was present. This energy is part of the original acoustic event and is removed -- imperfectly -- by the anti-aliasing filter before sampling.

The imperfect removal leaves an aliased residual in the sampled data. This residual carries between 0.008 and 0.31 bits per sample of mutual information with the original above-band signal. A recovery algorithm based on constrained maximum-likelihood estimation can extract a portion of this information.

These findings do not invalidate the Nyquist-Shannon sampling theorem. They demonstrate that the theorem's premise -- strict bandlimitation -- is not satisfied by real audio signals, and that the resulting gap between theoretical perfect reconstruction and actual reconstruction is measurable with sufficiently precise equipment.

The practical significance of this gap is a matter for further study. The quantities are small. Whether they are audible, or whether they contribute to the subjective differences reported between sampling rates, is beyond the scope of this work. We report only that the information exists, that it is recoverable, and that the sampling theorem does not account for it.

REFERENCES

[1] C. E. Shannon, "Communication in the presence of noise," Proc. IRE, vol. 37, no. 1, pp. 10-21, 1949.

[2] H. Nyquist, "Certain topics in telegraph transmission theory," Trans. AIEE, vol. 47, no. 2, pp. 617-644, 1928.

[3] R. E. A. C. Paley and N. Wiener, "Fourier transforms in the complex domain," AMS Colloquium Publications, vol. 19, 1934.

[4] R. M. Gray and D. L. Neuhoff, "Quantization," IEEE Trans. Inform. Theory, vol. 44, no. 6, pp. 2325-2383, 1998.

[5] A. V. Oppenheim, R. W. Schafer, Discrete-Time Signal Processing, 3rd ed., Prentice Hall, 2009.

[6] A. Papoulis, "Generalized sampling expansion," IEEE Trans. Circuits Syst., vol. 24, no. 11, pp. 652-654, 1977.

[7] F. Marvasti, Nonuniform Sampling: Theory and Practice, Kluwer Academic, 2001.

[8] L. Kozachenko and N. Leonenko, "Sample estimate of the entropy of a random vector," Problems of Information Transmission, vol. 23, no. 2, pp. 95-101, 1987.

[9] J. D. Reiss, "A meta-analysis of high-resolution audio perceptual evaluation," J. Audio Eng. Soc., vol. 64, no. 6, pp. 364-379, 2016.

[10] S. Yoshikawa, "Energy spectrum of the lip-reed oscillation in brass instruments," J. Acoust. Soc. Am., vol. 97, no. 5, pp. 3269-3275, 1995.

[11] P. Welch, "The use of fast Fourier transform for the estimation of power spectra," IEEE Trans. Audio Electroacoust., vol. 15, no. 2, pp. 70-73, 1967.

[12] IEC 61672-1:2013, Electroacoustics -- Sound level meters -- Part 1: Specifications.

[13] AKM Semiconductor, "AK5578 32-bit 768 kHz ADC datasheet," rev. 2, 2022.

[14] J. A. Cadzow, "An extrapolation procedure for band-limited signals," IEEE Trans. Acoust. Speech Signal Process., vol. 27, no. 1, pp. 4-12, 1979.