Real-time considerations for a source-time dominant auralization scheme

Andrew Christian
Structural Acoustics Branch
NASA Langley Research Center

Randall Ali
Dept. of Electrical Engineering (ESAT)
STADIUS Center for Dynamical Systems, Signal Processing and Data Analytics, KU Leuven

184th Meeting of the Acoustical Society of America, 1aCA
May 8th 2023, Chicago, IL

Auralization and Recording

- When one records a sound using a well-built hardware system (for instance, recording a plane flying overhead), there is no reason to worry about distortions other than those of the microphone.
  - Aliasing, imaging, the source characteristics of what we are recording, etc.
  - All of this necessarily happens in real time…
- In contrast, when making auralizations, one (may) have to pay attention to several aspects of the scenario being simulated:
  - Aliasing can arise if there is too-high frequency content for the receiver.
  - Images may creep in if the source is sampled at a very low rate.

→ How does this work? Can we come up with an auralization scheme that has the same robustness as hardware recording?

Aliasing in Auralization

- Setting the receiver sampling frequency too low can lead to reflections of the signal content off of the receiver Nyquist frequency.

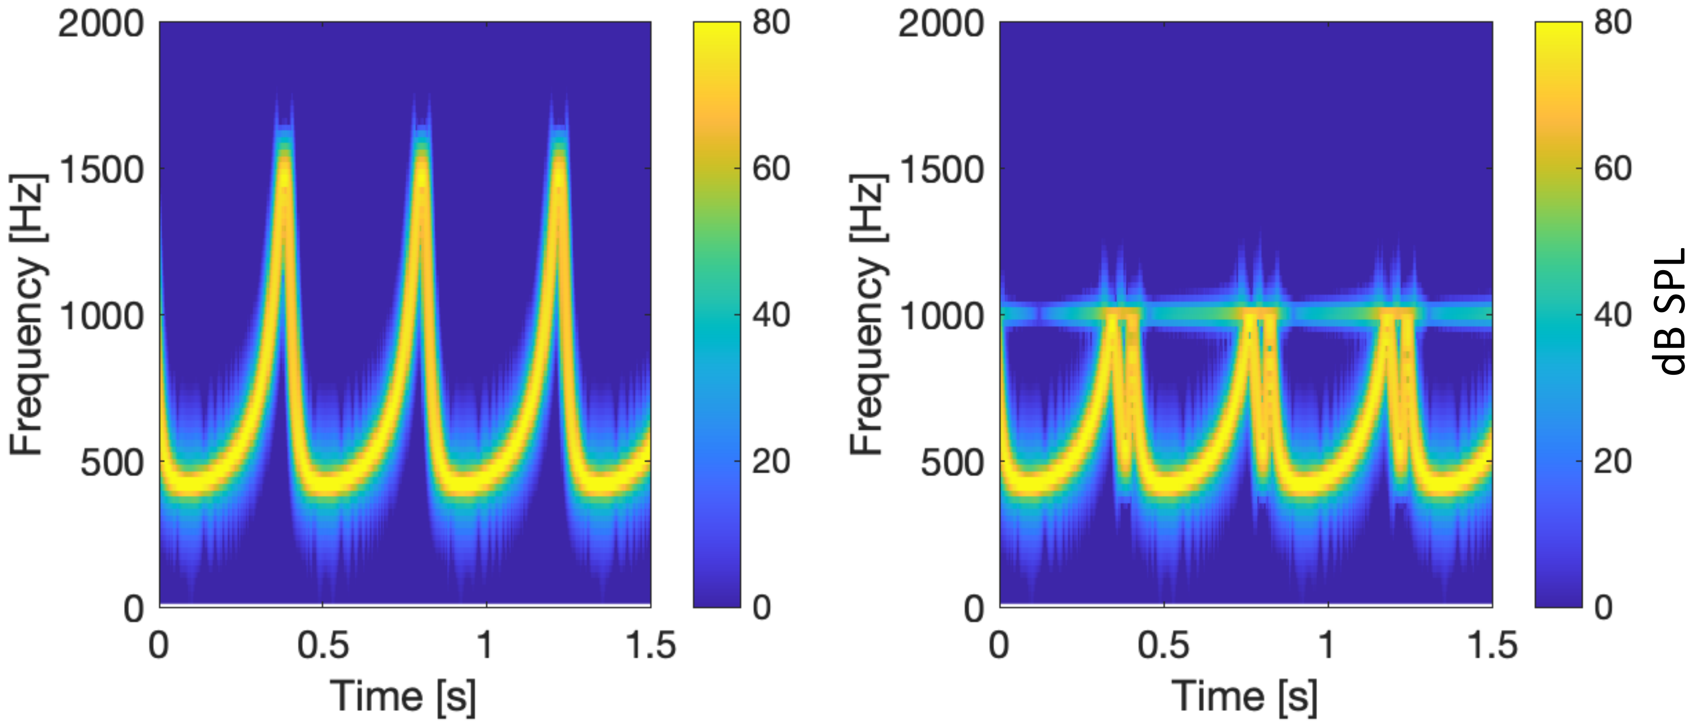

Imaging in Auralization

- If the source sampling frequency is very near to the source content, and the interpolation scheme used is not “clean,” images may be created that can propagate to the receiver.

- No interpolation is perfect…

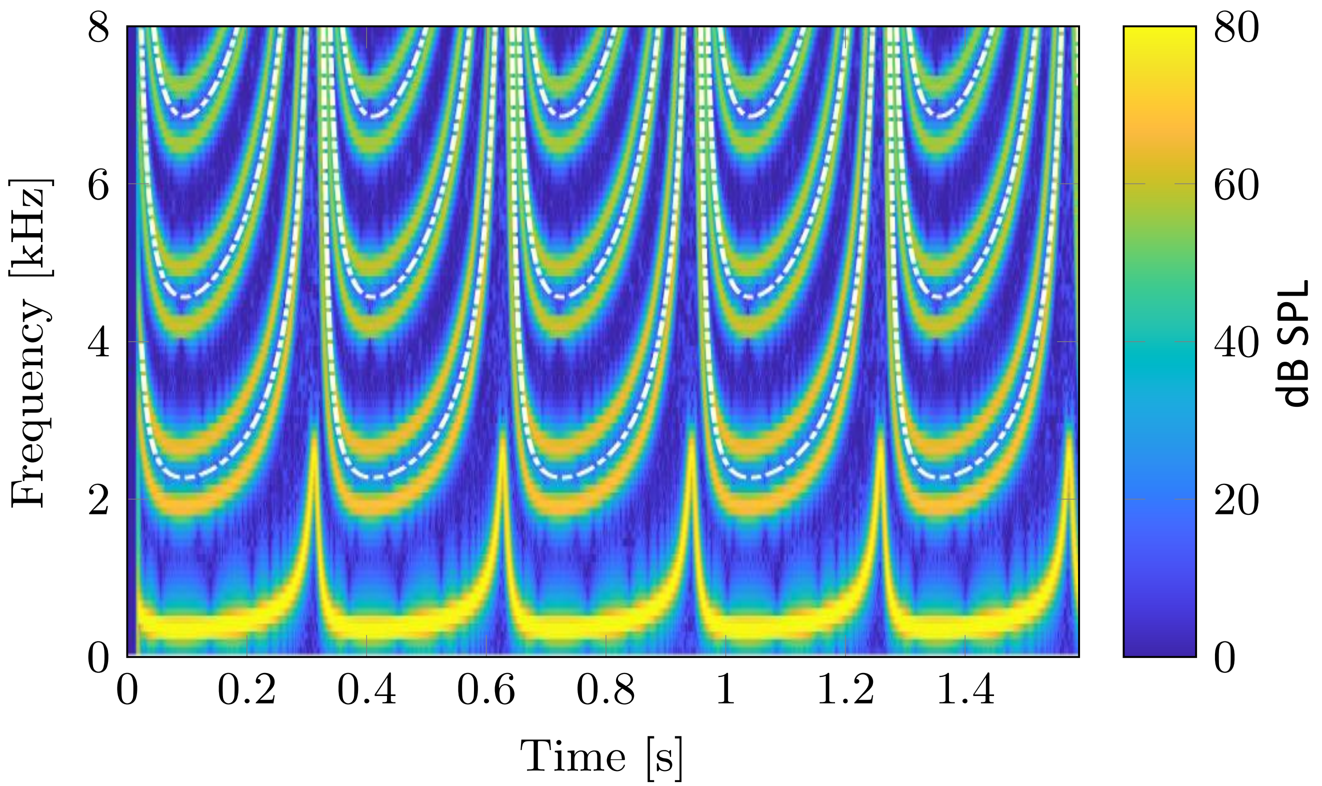

Imaging x Aliasing

- There is also, of course, the possibility of generating images that then become aliased…

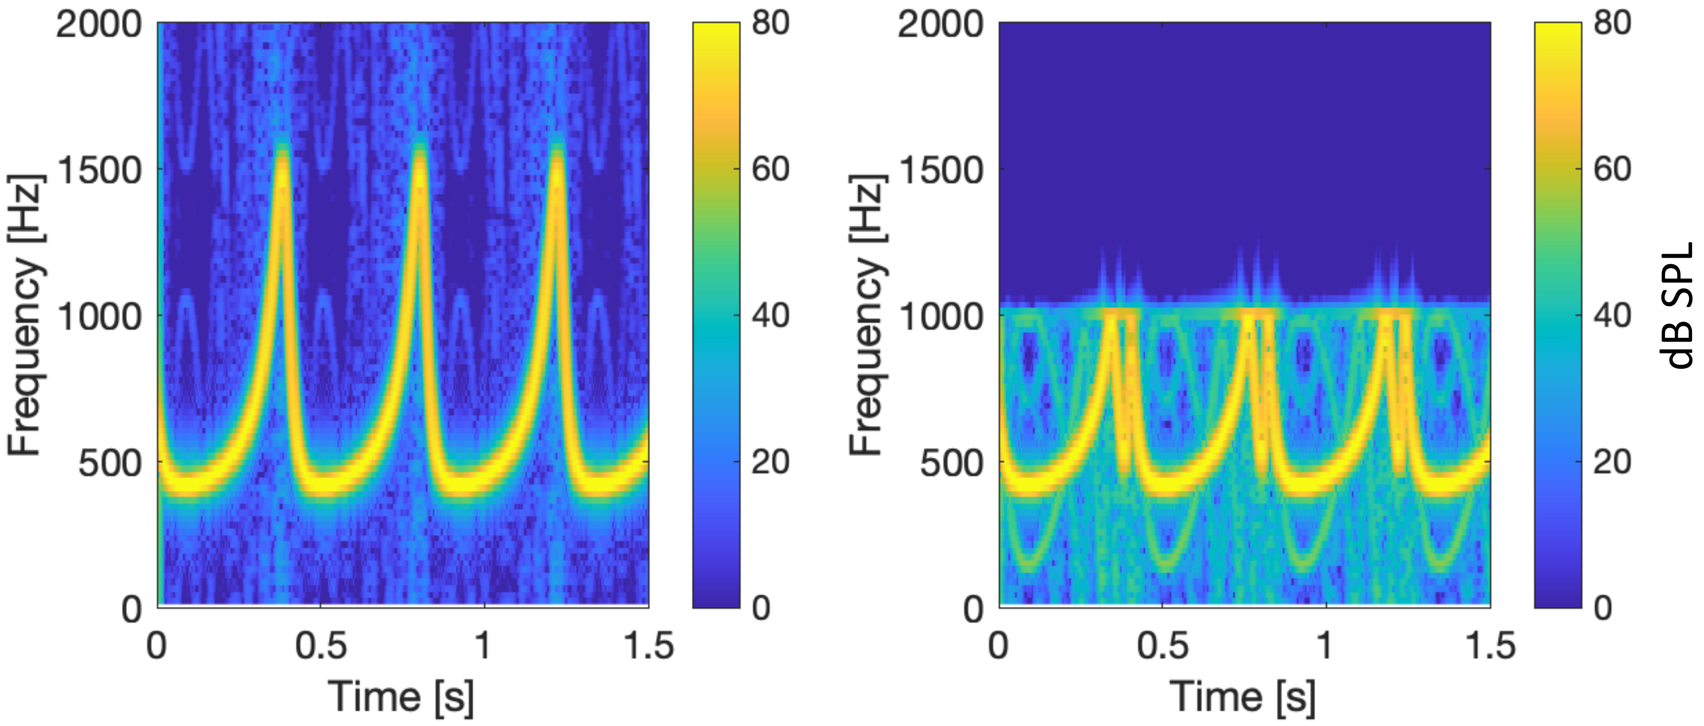

The NoTAP method

N-oversampled Transmitted Acoustic Pressure

- Define the desired receiver sampling frequency ($f_{s,R}$)
- Define a highly-oversampled domain to create a dense grid
- Interpolate the transmitted samples over this dense grid

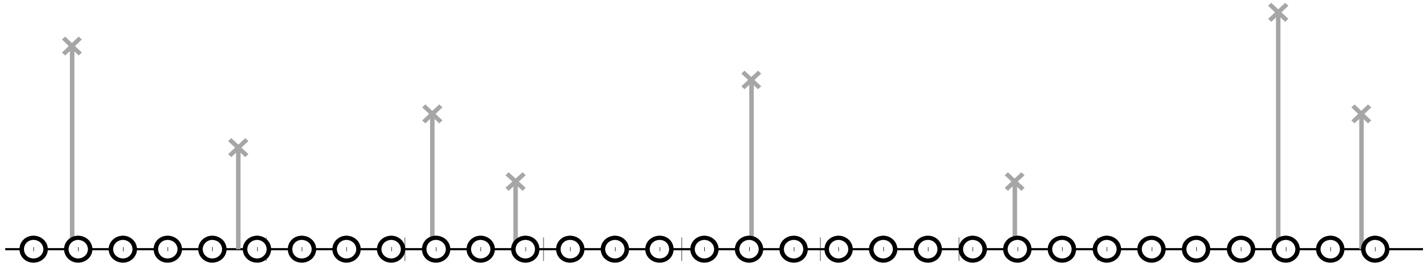

The NoTAP method

N-oversampled Transmitted Acoustic Pressure

Source → Transmitted → Oversampling + Interpolation

- We don’t have to use very complex interpolation methods
- Zero-order Hold → Copy the incoming sample to create a step function

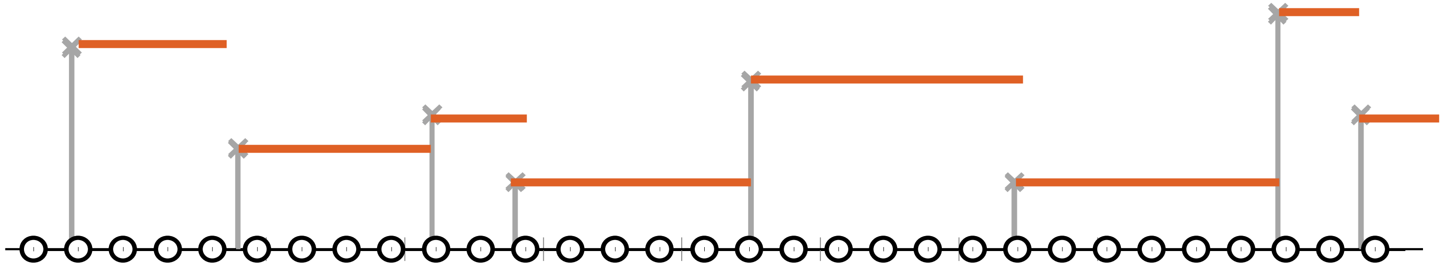

The NoTAP method

N-oversampled Transmitted Acoustic Pressure

Source → Transmitted → Oversampling + Interpolation → Low-pass + downsampling → Upsample to listening fs (optional)

- Lowpass filter and downsample to your desired receiver sampling frequency
    - This can be done with successive IIR decimations – a very fast process
    - This idea of oversampling, zero-order hold and then filter has been used for a long time in digital to analog conversion.

- The method keeps track of the rate at which the samples are arriving. If they arrive at a rate lower than $f_{s,R}$ then the lowpass filter can be used to eliminate frequency regions that could only be filled by images of the source content.

Applying NoTAP to the Signal

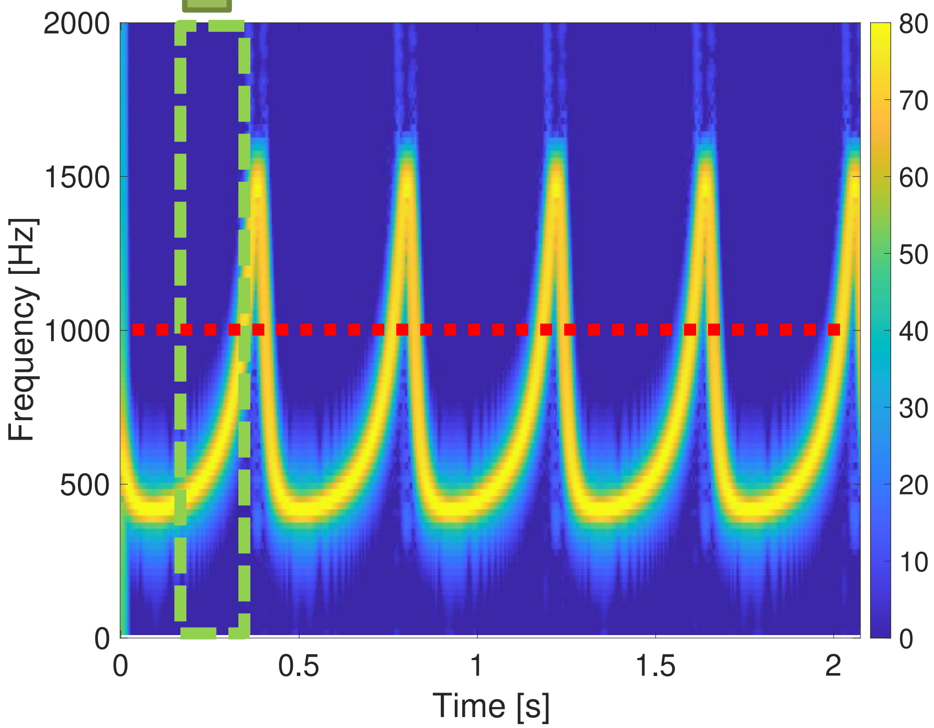

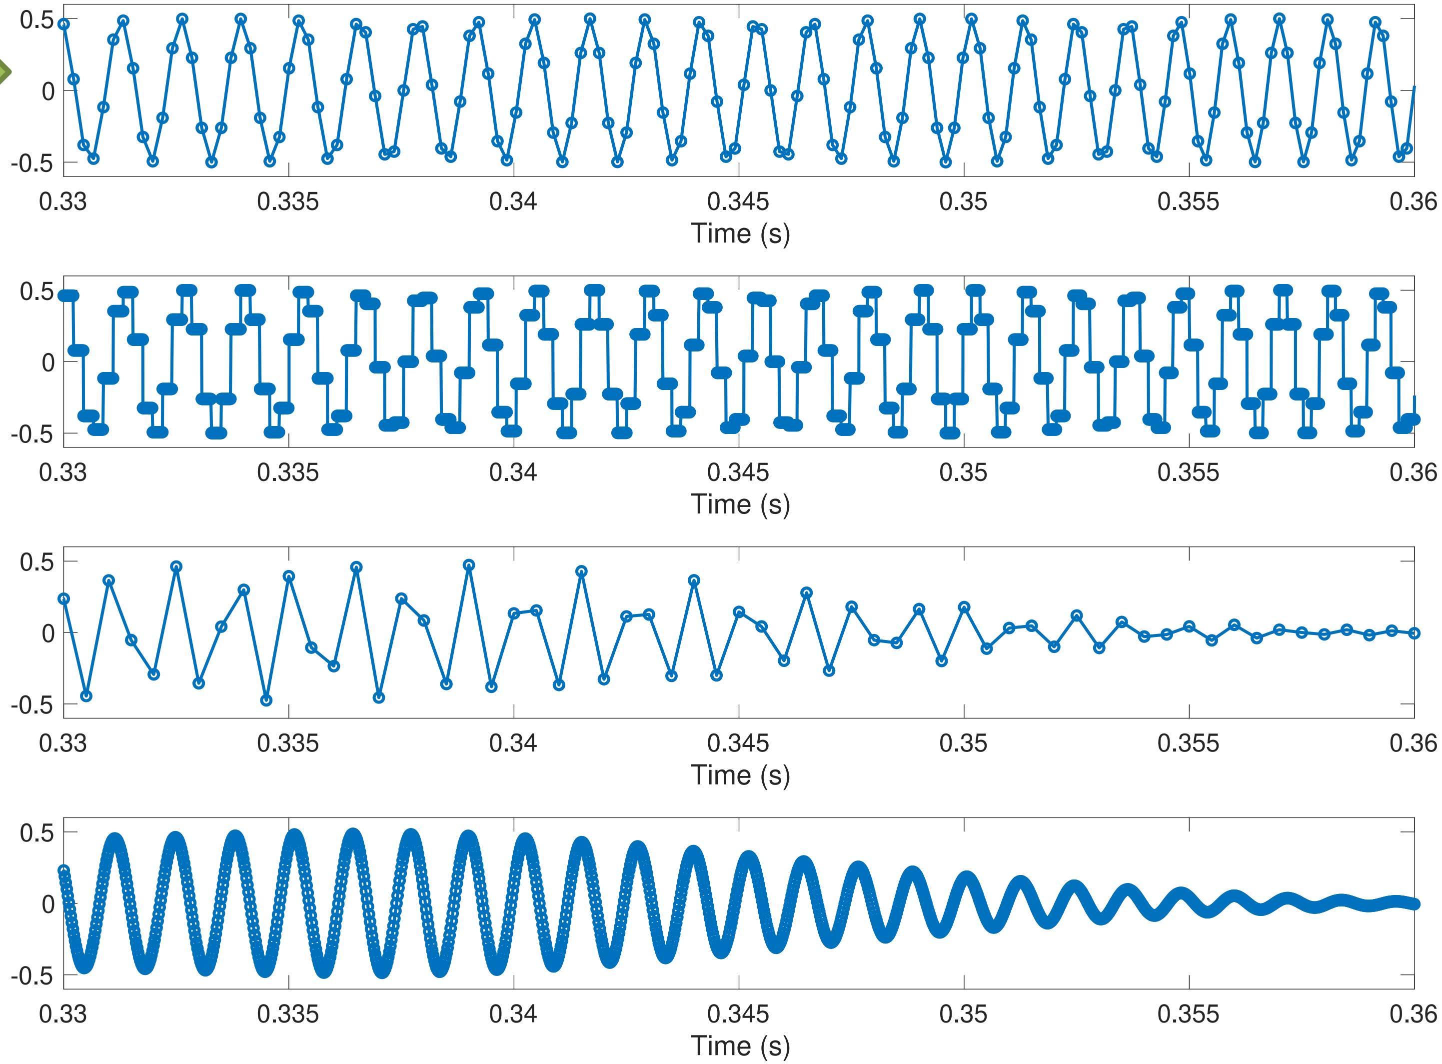

Real-Time Properties

- The NoTAP approach is particularly well-suited to real-time implementation.
- It is a source-time dominant scheme.
  - Samples are processed as they come from the source.
  - As opposed to receiver-time dominant (RTD) schemes where the receiver asks for samples from the source (two-way communication).
- It is a collection of basic DSP building blocks that can all run in synchrony with the final (output) sampling clock.
  - Copying information between samples.
  - IIR-based decimation.
- A very similar hardware-based scheme was proposed in the ‘80s.

The Catch: A Zero-Order Hold Noise Source

- When a new sample is received, the first sample to have that value will be AFTER the reception time.
- The "miss time" will be uniformly distributed over the oversampling intervals.
- The "miss magnitude" will be based on the difference between the previous received sample and the new one.

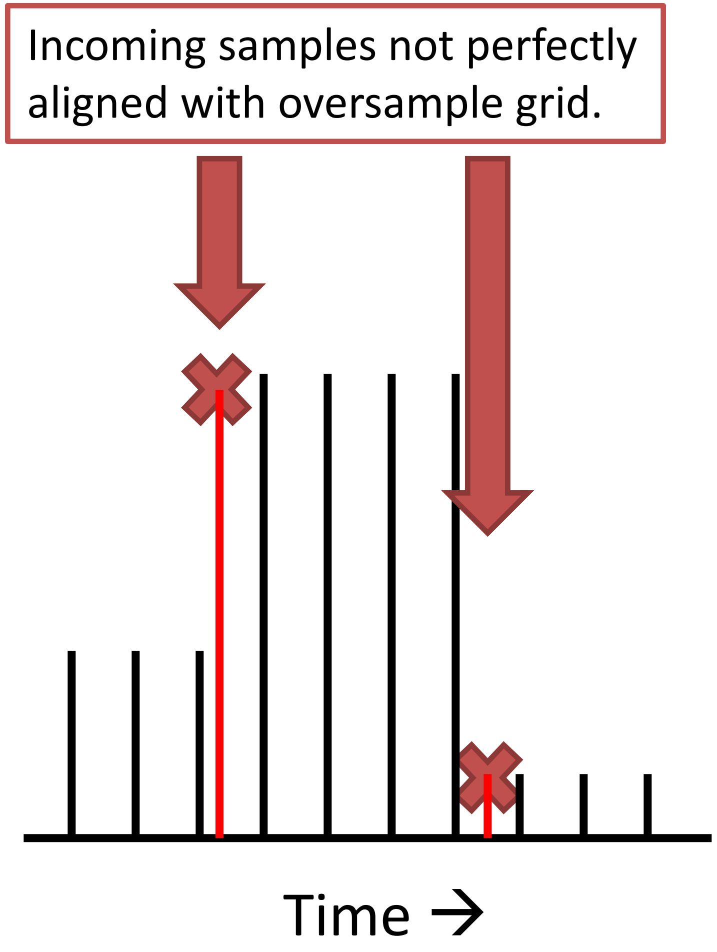

This is a white noise source with magnitude related to the oversampling ratio and the first derivative of the incoming signal.

The Catch: A Zero-Order Hold Noise Source

- When a new sample is received, the first sample to have that value will be AFTER the reception time.
- The "miss time" will be uniformly distributed over the oversampling intervals.
- The "miss magnitude" will be based on the difference between the previous received sample and the new one.

This is a white noise source with magnitude related to the oversampling ratio and the first derivative of the incoming signal.

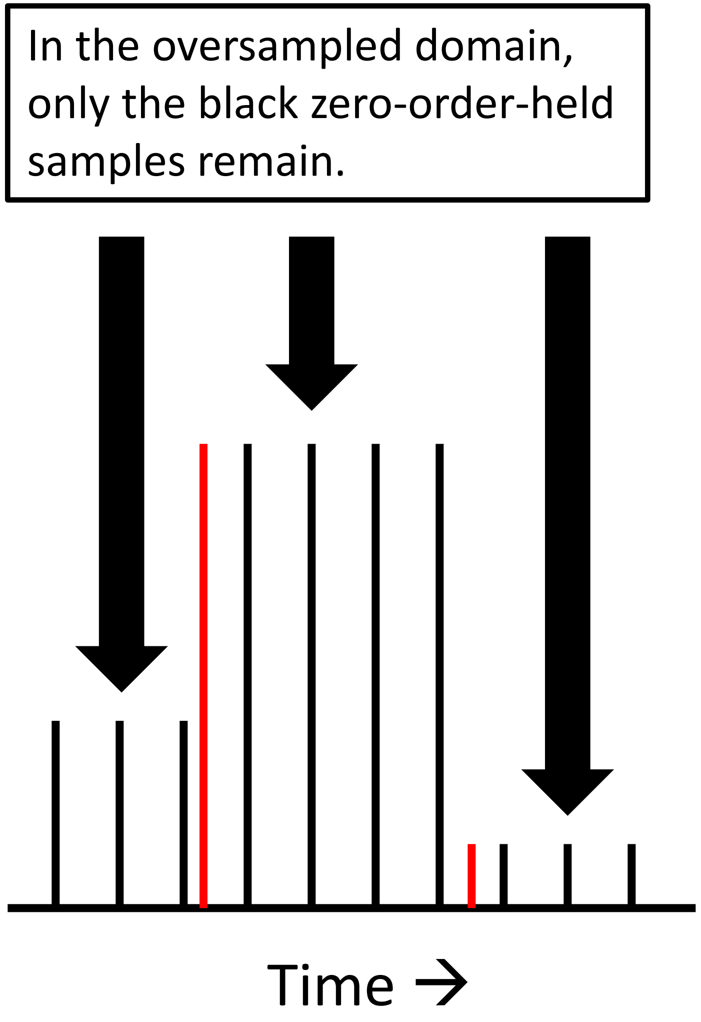

The Catch: A Zero-Order Hold Noise Source

- When a new sample is received, the first sample to have that value will be AFTER the reception time.
- The "miss time" will be uniformly distributed over the oversampling intervals.
- The "miss magnitude" will be based on the difference between the previous received sample and the new one.

This is a white noise source with magnitude related to the oversampling ratio and the first derivative of the incoming signal.

Interpolating from the black lines only gives the misses indicated by the dotted parts of the red line.

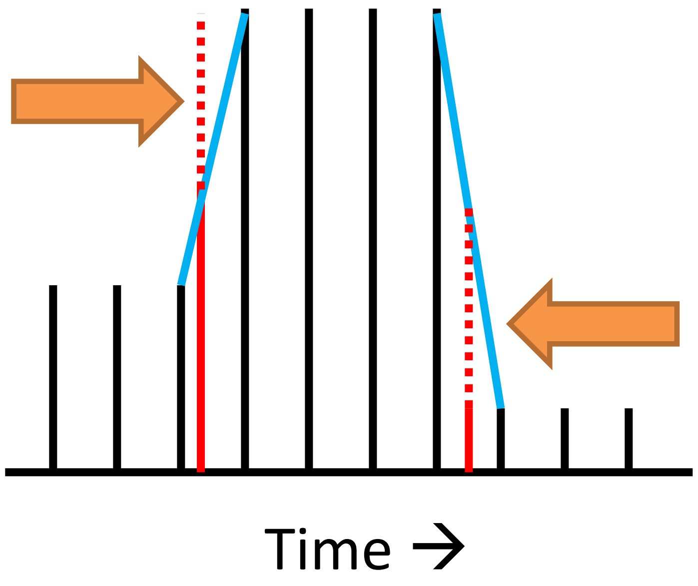

What Can Be Done?

- Staying with the ZOH approach, the oversampling ratio can be increased.
  - Increasing the ratio by an order of magnitude is really just a matter of copying a value 10 times and adding another decimation operation.
  - While a single increase of this type does not constitute a severe increase in computational burden, it certainly cannot be done indefinitely.
  - There will be a delay associated with the increased number of decimation steps.

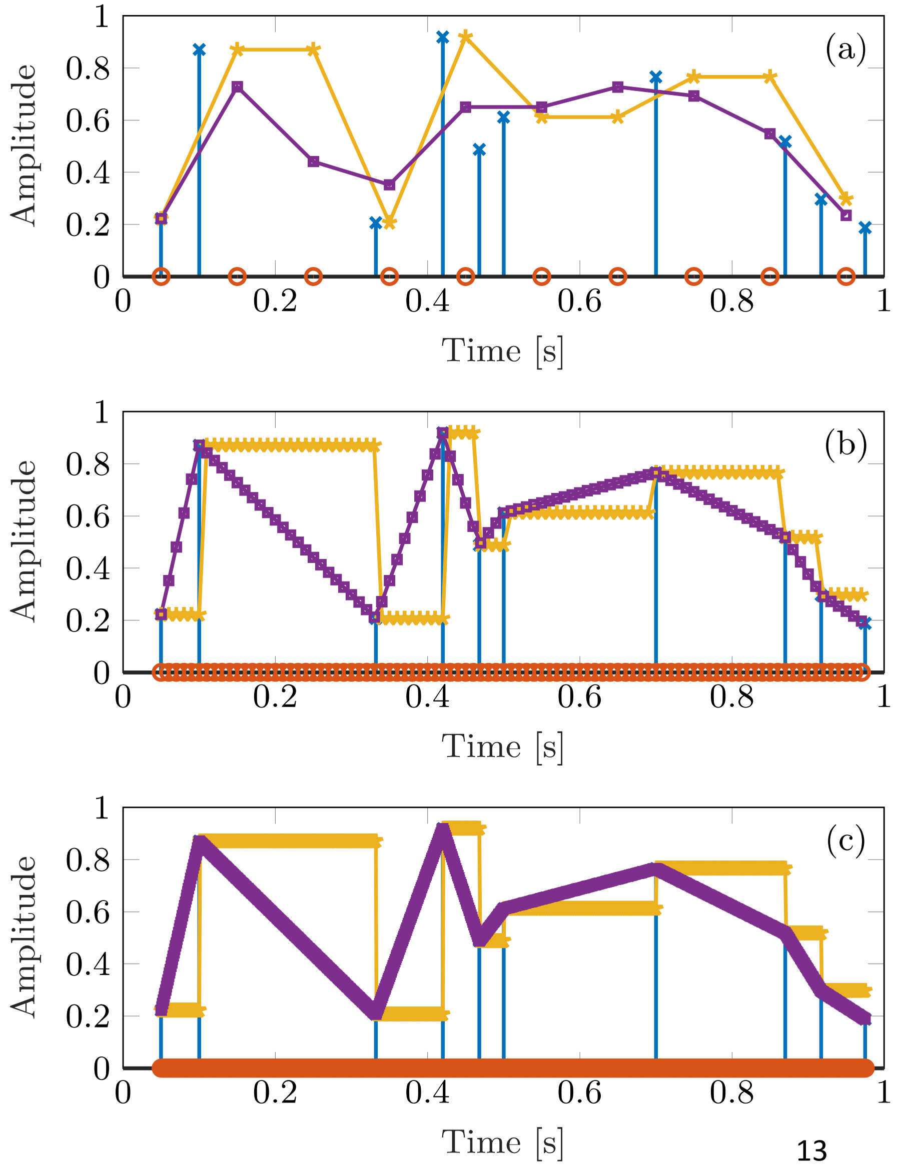

What Else Can Be Done?

- Higher-order interpolation can be used.
  - A “first-order hold” approach uses a linear interpolant between incoming samples.
  - Now not strictly real-time scheme – one must wait for the *next* sample in order to process points between two samples.
  - Provided that the incoming samples are stored with sufficiently high precision AND that their arrival times are stored as real values (i.e., not discretized to the oversampled grid), this noise source will be effectively eliminated.

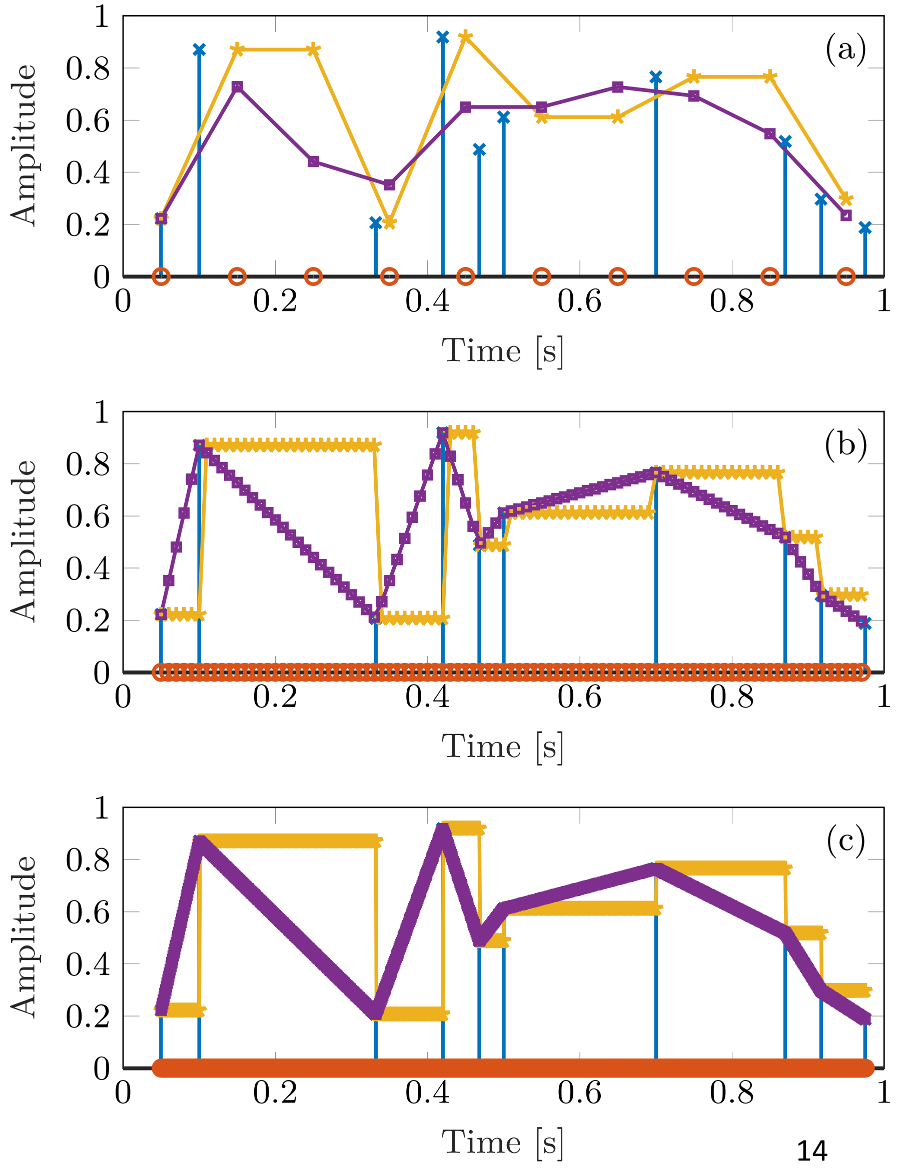

Effectiveness of These Approaches

- Changing these parameters of the approach (oversampling ratio, interpolation scheme) has the expected results.
  - Noise is greatly diminished with increase in oversampling ratio.
  - Noise disappears with linear interpolation.

RTD "Ground Truth"

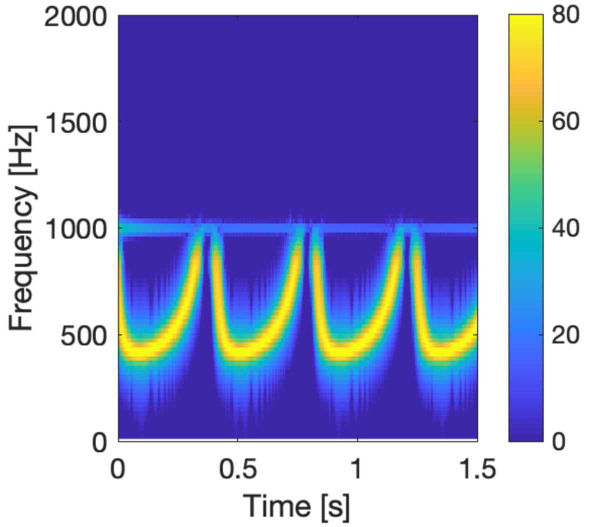

ZOH, 50x

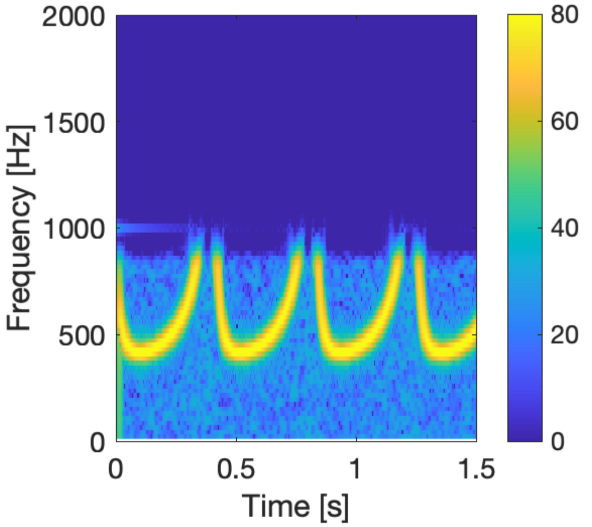

ZOH, 500x

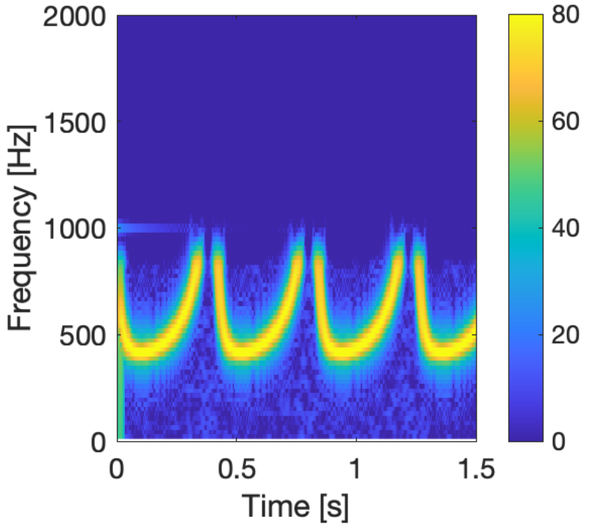

Linear, 50x

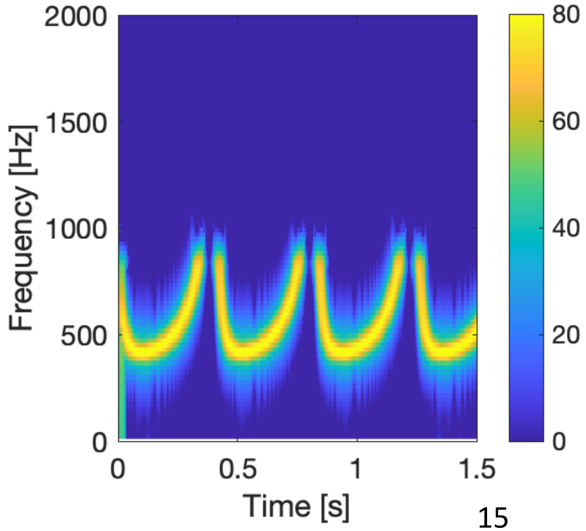

Conclusions

- The NoTAP method of source-time dominant auralization continues to be developed.
  - It addresses the possibility of both aliasing and imaging artifacts.
  - It can be implemented in an absolute-real time manner.

- In its real time form, a 3-way tradeoff exists between a source of white noise, the computational effort, and the nearness to true real-time performance.
  - Perhaps some more exotic solution to this problem can be cooked up in the future…

Questions?

Andrew Christian – andrew.christian@nasa.gov

Randall Ali – randall.ali@esat.kuleuven.be