

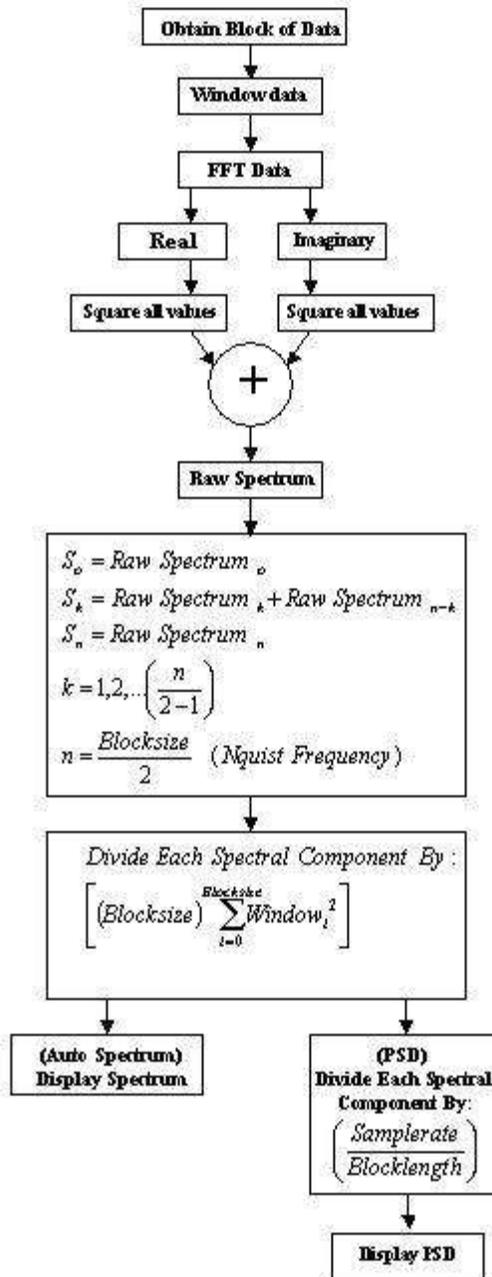
IADS Frequency Analysis FAQ

(Updated: March 2009)

** Note - This Document references two data set archives that have been uploaded to the IADS Google group available in the "Files" area called; IADS Frequency ATP and IADS Frequency FAQ.*

1) Can the math process for Auto Spectrum and Power Spectral Density be described in detail?

Answer – Below is a high level flow chart on the math process which can also be found in the IADS Help System on the "Power Spectral Density" page. It shows how the Auto Spectrum and PSD are related. Other topics of interest are window equivalent noise bandwidth and peak amplitude, both of which scale the amplitude of the Auto Spectrum. A comparison of results with another application using known inputs might be helpful as well.

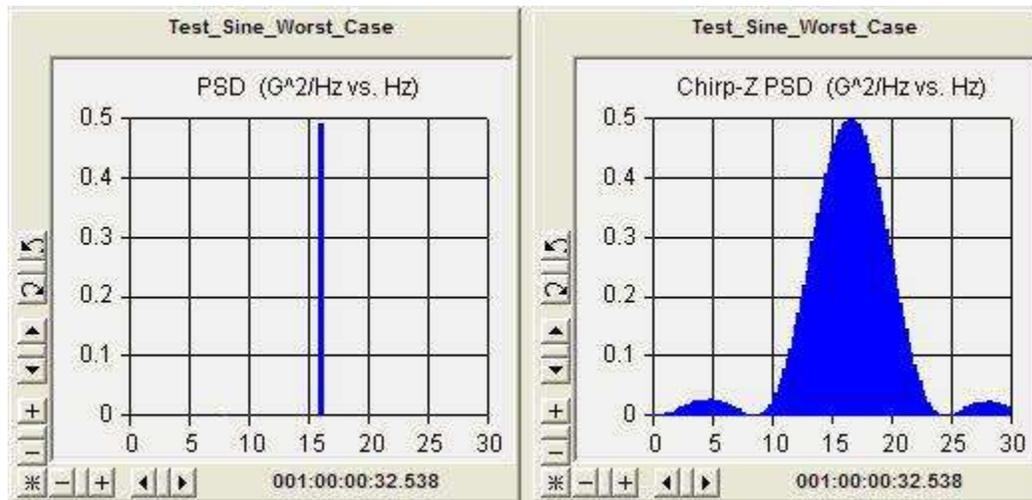


2) Can you describe the different characteristics between Auto Spectrum and Power Spectral Density for frequency plots?

Answer - The main difference in the FFT-based PSD and “un-scaled” Auto Spectrum is a (samplerate/blocksize) factor. There are certain scalings in the Auto Spectrum that have been added starting with version 6.0 such as window equivalent noise bandwidth and peak, that also add to potential differences.

3) Can you describe the different characteristics between CZT and FFT?

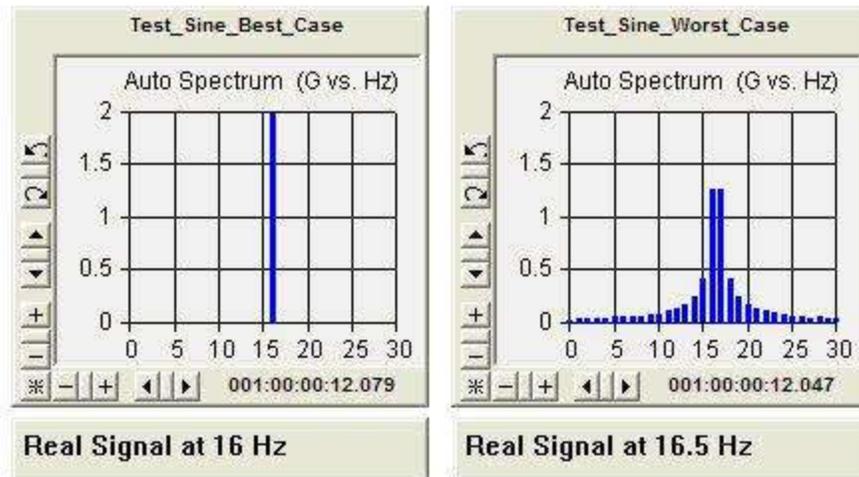
Answer - The FFT distributes spectral energy between 0 Hz and the Nyquist Frequency at regular intervals. The CZT distributes spectral energy between the lower and upper display limits at regular intervals.



The CZT is useful for obtaining greater resolution in a narrow frequency band, and you may find that HPD works better with CZT. The CZT is also more computationally complex.

4) Can you define frequency plot leakage problem, and how its effects can be determined?

Answer - Spectral leakage is a well documented phenomenon (http://en.wikipedia.org/wiki/Spectral_leakage), and it happens in our Frequency Plot because we use FFTs. Windows such as Hanning, and Hamming are used to minimize leakage. The effects of leakage can easily be seen in the tables in Item 5. The energy in the frequency bin (the peak) is leaked into adjacent bins.



5) By changing the parameter's sample rate, the frequency plots output would change magnitude. Can you explain why this happens and how it can be minimized? This was seen when using and modifying the IADS Frequency ATP data.

Compared a derived 5k sps and 2k sps with original 1024 10 hz 1.414 Peak to peak sine wave both Auto Spectrum and PSD.

1. **Block size= 4096 the 5k sps peak equals .764**
2. **Block size= 2048 the 2k sps peak equals .844**
3. **Block Size=1024 the Original 1024 sps peak equals 1.0**

Peaks outputs all equaled 1.0 after converting sps.

1. **Block size= 4096, 5k sps to 4096 sps peak equals 1.0**
2. **Block size= 2048, 5k sps to 2048 sps peak equals 1.0**
3. **Block Size=1024 the Original 1024 sps peak equals 1.0**

*Answer – Using IADS versions prior to 6.0, a 1.414 G Peak Sine (2.828 Peak to Peak) will give you 1.0 G**2 RMS in the Auto Spectrum. If the samplerate/blocksize is equal to 1, the PSD and AutoSpectrum will give you the same amplitude with different units. Version 6.0 gives you the ability to view Auto Spectrum amplitudes in “unsquared” peak instead of RMS.*

A change in sample rate while holding the block size constant greatly affects the amplitude shown on a PSD. If you double the sample rate, the Nyquist frequency also doubles, but the number of frequency bins remains the same. Each frequency bin is effectively twice as wide and half as tall (the area under the PSD is the power).

There is a term called “scalloping loss” that describes the ripple in the amplitude of an FFT (Online explanations of scalloping are more succinct than I'll ever be (http://en.wikipedia.org/wiki/Window_function#Discrete-time_signals)). It may happen

that the sample rate will be an integral multiple of the frequency of interest and you'll get the full magnitude. If not, energy will leak and the apparent peak will not be the true frequency. Using the Flat Top Window (available since version 5.3.3) will give you surprisingly good amplitude numbers even at worst case (signal frequency exactly between two frequency bins).

6) Can you define how averaging frequency, #Blocks, and %Overlap effect data within frequency plots?

Used the IADS Frequency ATP data.

- i. Assigned FivetoTwoFiftyFiveLogSweep parameters to 2 frequency plots. One plot having frequency averaging off and the other plot having frequency averaging on, #Blocks=0, and %OverLap=10. The plots look completely different.
- ii. Same setup as above, but with second plot #block=1. The plots look very similar.
- iii. Same setup as above, but with second plot #block=2. The second plot peak was reduced by at least half, but bandwidth increased.

Answer – “Frequency” averaging refers to the averaging of data at the same frequency over multiple FFTs (as opposed to 3 or 7 point “Smoothing”). The averaging is turned on when you choose “Frequency” and off when you chose “None”.

Averaging tends to dilute dynamic peaks over time - if you have a transient, the energy in that transient is caught by the FFT Block encompassing that time, so each block (in the average) before and after it will diminish the amplitude shown. Higher overlap percentages increase the likelihood that a transient will contribute energy to more than one FFT. Also, since less time is required between FFTs with a higher overlap, screen updates happen more often.

If you're looking at sinusoids, averaging and overlap shouldn't have much influence on the displayed data. Averaging over a sweep will do exactly what you said.

7) Does %OverLap work even if frequency averaging is NONE?

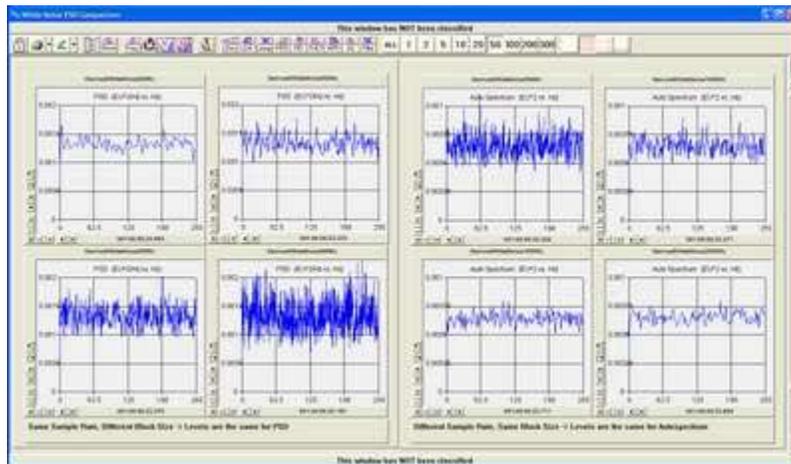
Answer – Yes, screen updates happen more often with higher overlap percentages.

8) Can you define Windowing effects for all windows defined in frequency plots? Windowing types: Rectangular, Hanning, Hamming, Blackman, and Kaiser-Bessel, (with alpha).

Answer - Yes, but there are more rigorous treatments out on the internet (http://en.wikipedia.org/wiki/Window_function#Discrete-time_signals). I'd recommend giving our Flat Top Window a try. Some of your concerns may be lessened once you see what that window can do.

9) Auto Spectrum and PSD amplitudes will often vary with block size. On flight data, amplitude has increased with block size up to a point and then will lower again. Does it look like we're achieving an optimum block size value? We are using 500 sps with an 8192 block size, on our current high modal density large transport aircraft.

Answer – Following is the effects of white noise:



1. Same sample rate/different block size gives same levels on a PSD
2. Same blocksize/different sample rate gives same levels on an Auto Spectrum.

The lower amplitudes you are seeing are most likely caused by leakage (8192/500 is not a round number and I don't know if you are referring to a sinusoid or broadband signal here). I think that your method approaches a favorable scalloping level at some block size.

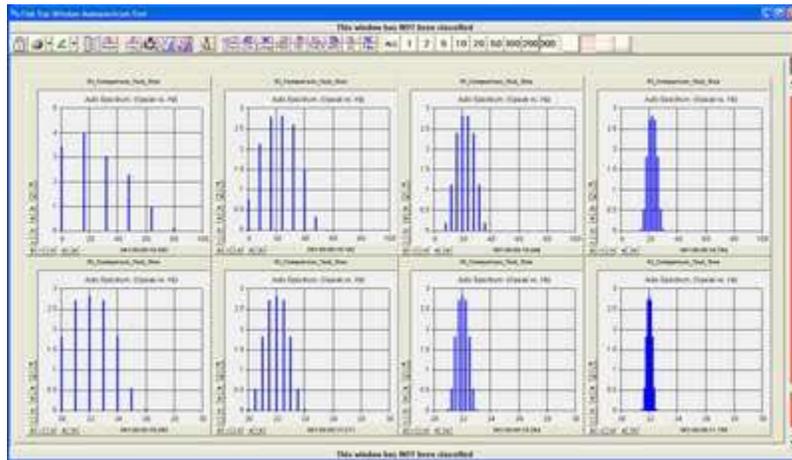
Please remember to treat the Auto Spectrum and PSD differently - how block size and sample rate affect them is different for different types of signals. Again, please try our Flat Top Window. You may find that the relative amplitude of modes is more important than absolute value (modal is more about frequency and damping).

10) Can you describe the characteristics on how changing the block size on frequency plots can change the amplitude of Auto Spectrum, PSD of both FFT and CZT data?

Answer – See Item 9 for a description of how the PSD and Auto Spectrum respond to block size changes with a white noise signal.

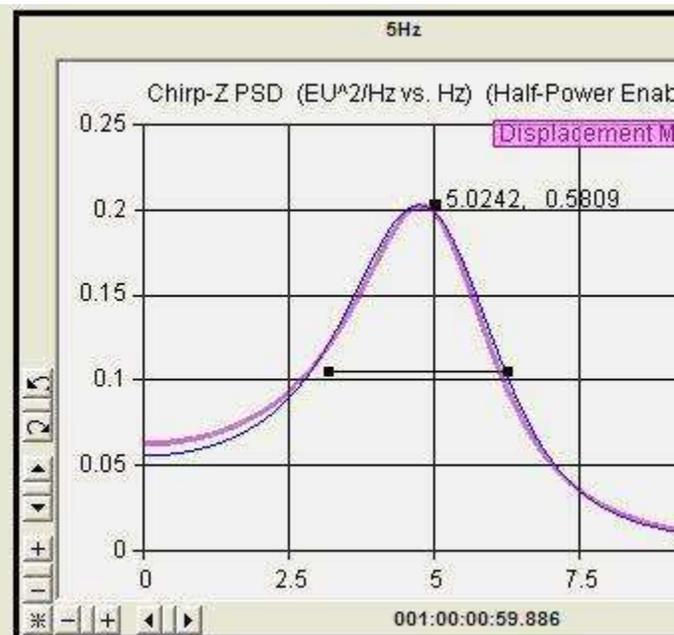
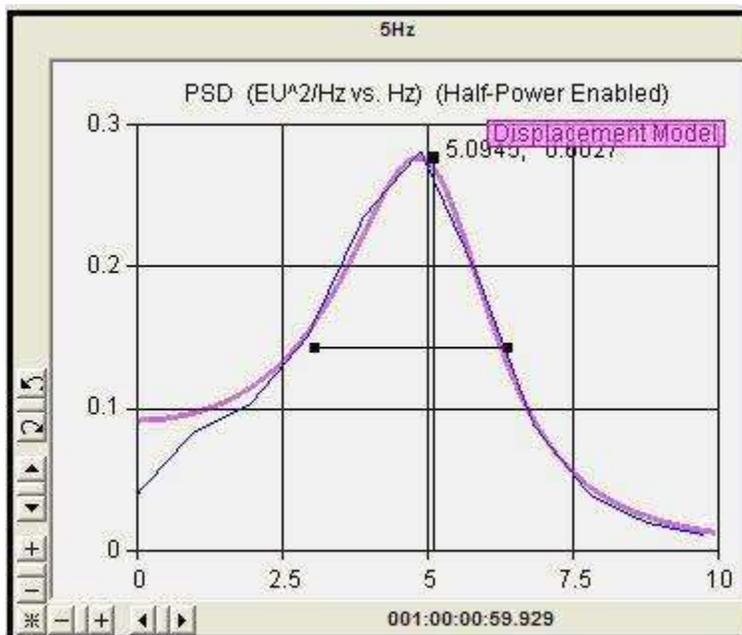
The IADS Frequency FAQ data set has IADS Desktops showing FFT-based comparisons of block size and sample rate. CZT is a little sketchy because of the possible variability of the upper and lower display limits.

These displays are using the Flat Top Window.



11) On Half Power Damping we usually adjust the overlay horizontal bars to achieve what we think is an optimum fit. And this would be to have an almost equal amount of area outside of the overlay as there is inside the overlay (since the waveform and overlay borders rarely match up).

This process greatly affects the damping estimate and seems more of an art than science. Are we doing this correctly? Are there better rules for fitting?



Answer - It is a damping "Estimate". I'd recommend comparing IADS results with your own model. It has been our experience that the FTE's already have an idea what the answer should be and they are using HPD to quickly verify/disclaim their assumption. If you are looking at the data for the first time, using the HPD as one of many tools is a good approach.

You may find that using the CZT will give you an easier fit because more points in the curve are there around the peak. The CZT also can reveal other peaks that are obscured by the FFT. Both of them are really samples in the frequency domain of a continuous spectrum.

The overlay is supposedly an ideal curve. Please try the "O" key hit. The nature of the signal (like displacement or acceleration) changes the way the curve rolls off. Also, with high damping, the damped natural frequency and the un-damped natural frequency start to diverge. If you drag the peak of the HPD to the right slightly you may get better results. Leakage can appear to add to damping with very lightly damped modes (measure an un-damped sine that has had a window used).

12) Can you define best practices from other flight researchers that use IADS frequency tools for either acoustic or flutter data analysis?

Answer - We don't keep that type of information at IADS, mainly because we build the tools at the request of our FTE users. However, I can try and explain how the tools work and suggest how they might be used to satisfy a requirement or avoid a dangerous flight condition if more questions are asked.

13) How could IADS be used to handle multiple frequencies? For example, can the Sum Area RMS function be used on several (say 4) peaks simultaneously?

Answer - Not individually like you said. If I understand what you're asking, we'd need 4 pairs of red bars and a multi-parameter SARMS Strip chart. You could run four SARMS PSDs on the same parameter.

The RFP Curve Fit in the Frequency Response Plot can handle up to 10 modes (frequency and damping) in real-time. The Time History Curve Fit can handle 10 modes as well, with frequency, damping, amplitude and phase.

The Sum Area RMS is a good tool for "un-leaking" a single peak.