The Effect of Time-Domain Interpolation on Response Spectral Calculations

David M. Boore

This note confirms Norm Abrahamson's finding that the straight line interpolation between sampled points used in the Nigam and Jennings (NJ) response spectral computation can lead to differences in response spectra compared with spectra computed from records for which the time series are interpolated properly, using the sampling theorem (which amounts to convolution in the time domain with a sinc function; see, e.g., Press *et al.*, 1992, p. 495). The interpolation enters into the NJ computations in two ways: 1) the standard NJ method uses analytic equations for the oscillator response, assuming linear acceleration between sampled points; and 2) the extended NJ method resamples the time series if the oscillator period is less than some multiple of the sampling interval (usually a factor of 10); the resampling assumes linear acceleration between the sample points. The differences are particularly likely to arise when the Fourier acceleration spectrum (FAS) has little decay with increasing frequency as the Nyquist frequency ($sps/2$, where sps is the samples per second of the digital record) is approached (although as the example here shows, the differences can arise even time series for which an antialiasing filter has been used). Such records are likely to be encountered on hard rock sites when the sample rate of the digital record is low.

Norm Abrahamson has processed many records, primarily from eastern North America, and has prepared various plots, including ratios of response spectra computed using the two interpolation assumptions vs. frequency. His plots show clear differences in the response spectra. My intention here is not to duplicate his work (an impossible task anyway), but simply to document the small amount of work that I have done, in the hope that others might find it useful. Norm deserves full credit for finding and investigating the issue; let's hope that he publishes his results.

I found a suitable record to investigate, sent to me by one of the NGA-E participants (Albert Kottke? Christine Goulet? Chris Cramer?). The three component files have the names *US.NCB.BHE.cor.acc*, *US.NCB.BHN.cor.acc*, and *US.NCB.BHZ.cor.acc*. The sensor output was sampled at 40 *sps*. The headers in the files indicate that they were not filtered. I converted the files to smc format and used my TSPP programs to compute FAS, response spectra, etc. Here are the FAS for the original time series (not interpolated):

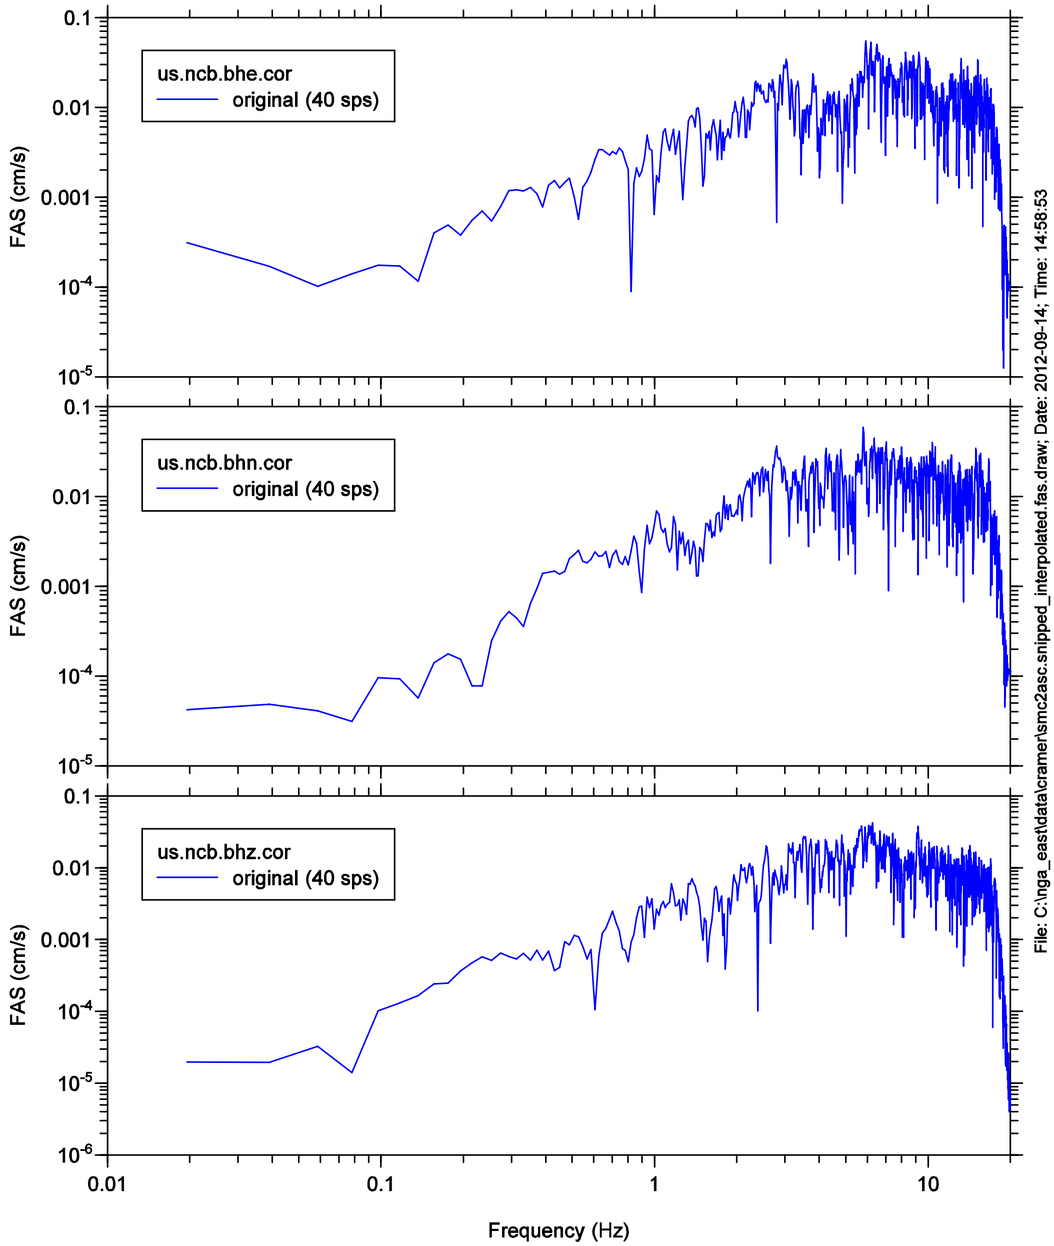

Figure 1. FAS of the original (non-interpolated) records.

Note that the FAS decay slowly, if at all, with increasing frequency until about 17 Hz, at which point they decay rapidly. This presumably is due to an antialiasing filter included at some point in the data stream. Because the Nyquist frequency is 20 Hz, the anti-aliasing seems quite sharp (with a transition band between about 17 and 20 Hz), which might lead to some of the effects shown later.

I interpolated the time series to 160 *sps* in the following standard way (also used by Norm): I snipped out a 50 s portion of the record containing the signal, extended the original time series to a power of 2 (from 2000 to 2048 samples, in this case) by padding with zeros, computed the FAS using a FFT, extended the spectrum with zeros beyond the Nyquist frequency, up to a "new" Nyquist frequency given by some multiple of two of the original Nyquist frequency (from 20 to 80 Hz, in this case), and then transformed back to the time domain. Because the frequency spacing of the FAS was not changed, the procedure resulted in a time series sampled at 160 sps, rather than the original 40 *sps*. The program for doing the resampling is *smc_interpolate_time_series_using_fft.for*, available in my TSPP package of time-series processing programs (see the online software link on www.daveboore.com). (I'll eventually add an option in my response spectrum main programs *smc2rs* and *smc2rs2* to include the resampling as an option).

A portion of the original and resampled time series is shown in Figure 2.

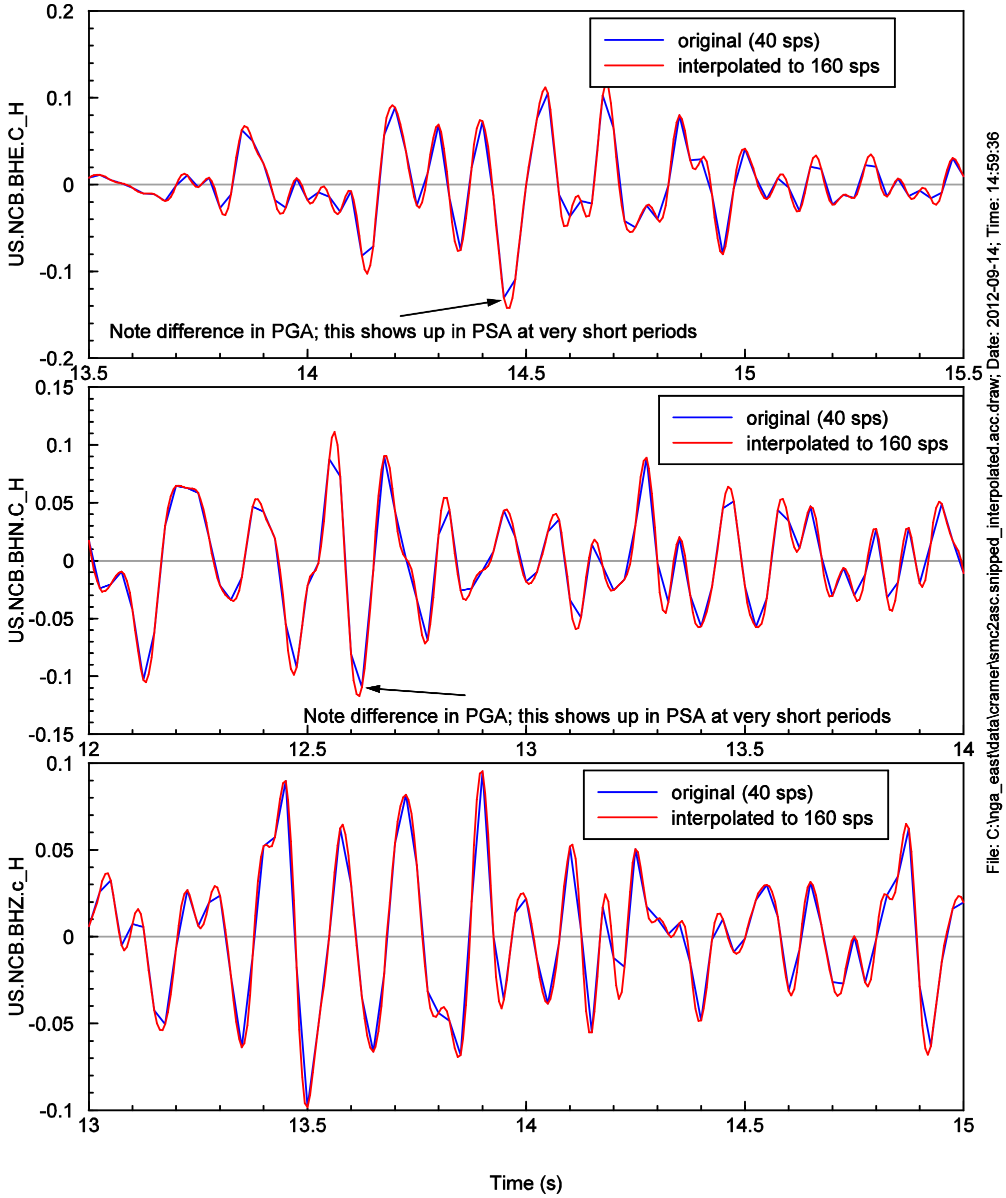

Figure 2. The portions of the records containing the absolute peak motions, comparing the straight line and sampling theorem interpolations.

Note that the resampled time series smoothes the vertices in the original time series, but in doing so occasionally introduces some extra oscillations. The absolute peak amplitudes (PGA) are generally increased slightly; this will result in an increase in the response spectra at short periods, as shown in the next figure.

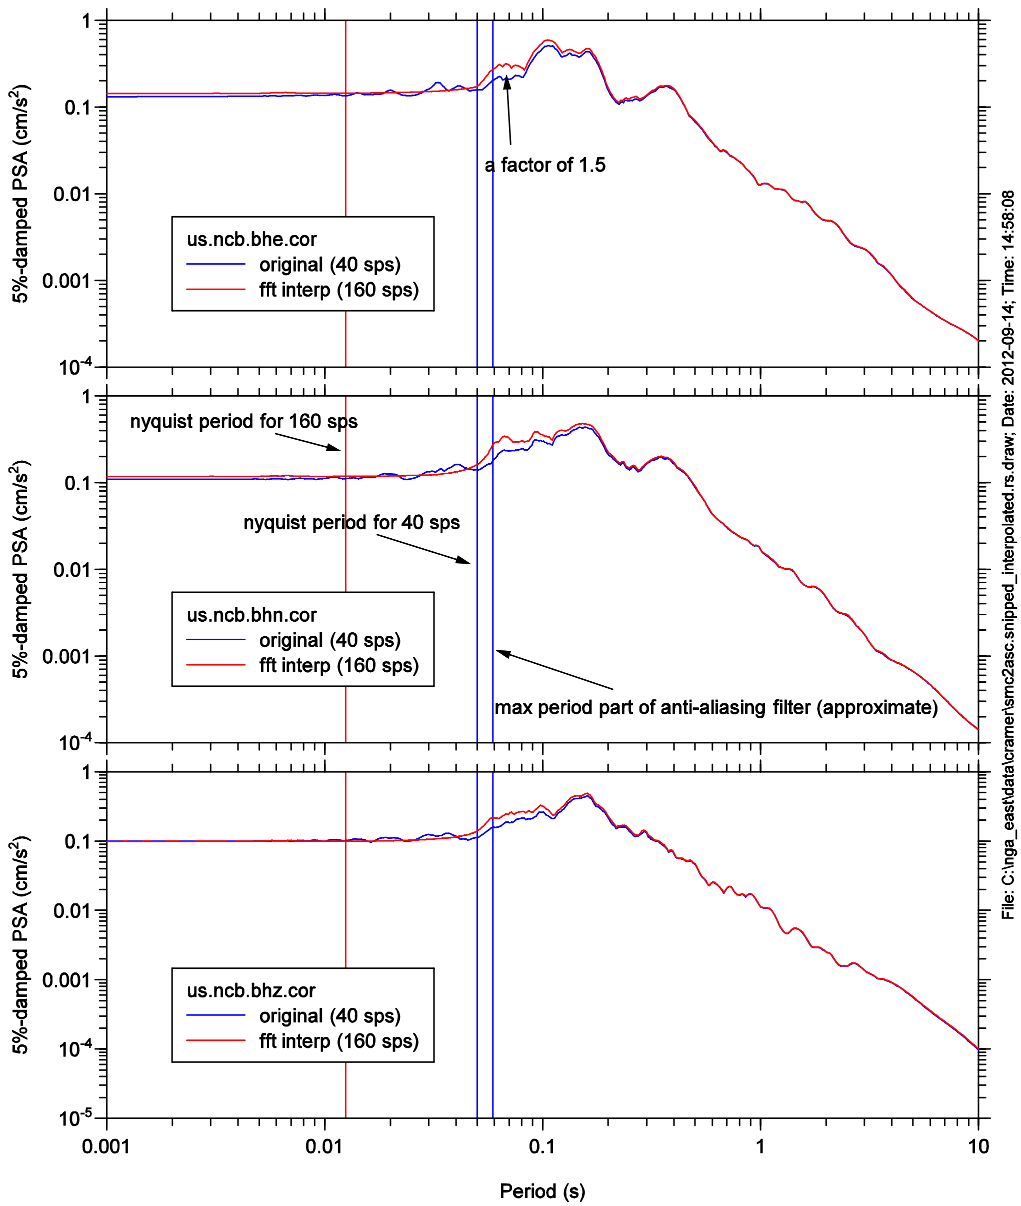

Figure 3. 5%-damped PSA for the original and resampled time series. For reference, the vertical lines show the periods corresponding to the Nyquist frequencies of the original and resampled time series. Also shown is the period corresponding to the frequency at which I estimate the anti-aliasing filter started (about 17 Hz).

There are significant differences in the PSA for periods exceeding the period corresponding to the start of the anti-aliasing filter (e.g., a factor of about 1.5 at a period of 0.08 s (12.5 Hz)). As seen in Figure 1, the FAS is relatively level for frequencies near 12.5 Hz (and as will be shown in the next figure, the FAS for the original and time series resampled to 160 *sps* are the same for frequencies less than the Nyquist frequency of 20 Hz).

To simulate a less severe anti-aliasing filter, I tapered the FAS of the original time series from 10 to 20 Hz before resampling. Figure 4 compares the FAS for the original, resampled to 160 *sps*, and tapered and resampled to 160 *sps* time series.

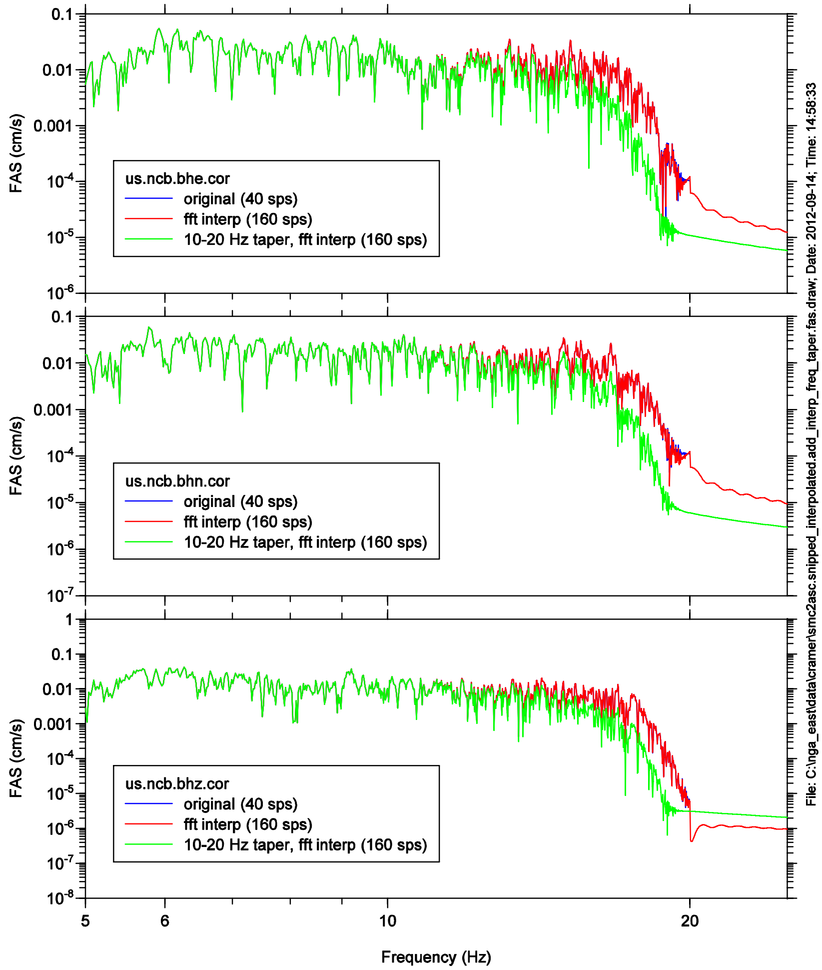

Figure 4. FAS of original and resampled time series, without and with an additional taper from 10 to 20 Hz.

A small portion of the three time series surrounding the peak motion is shown in Figure 5.

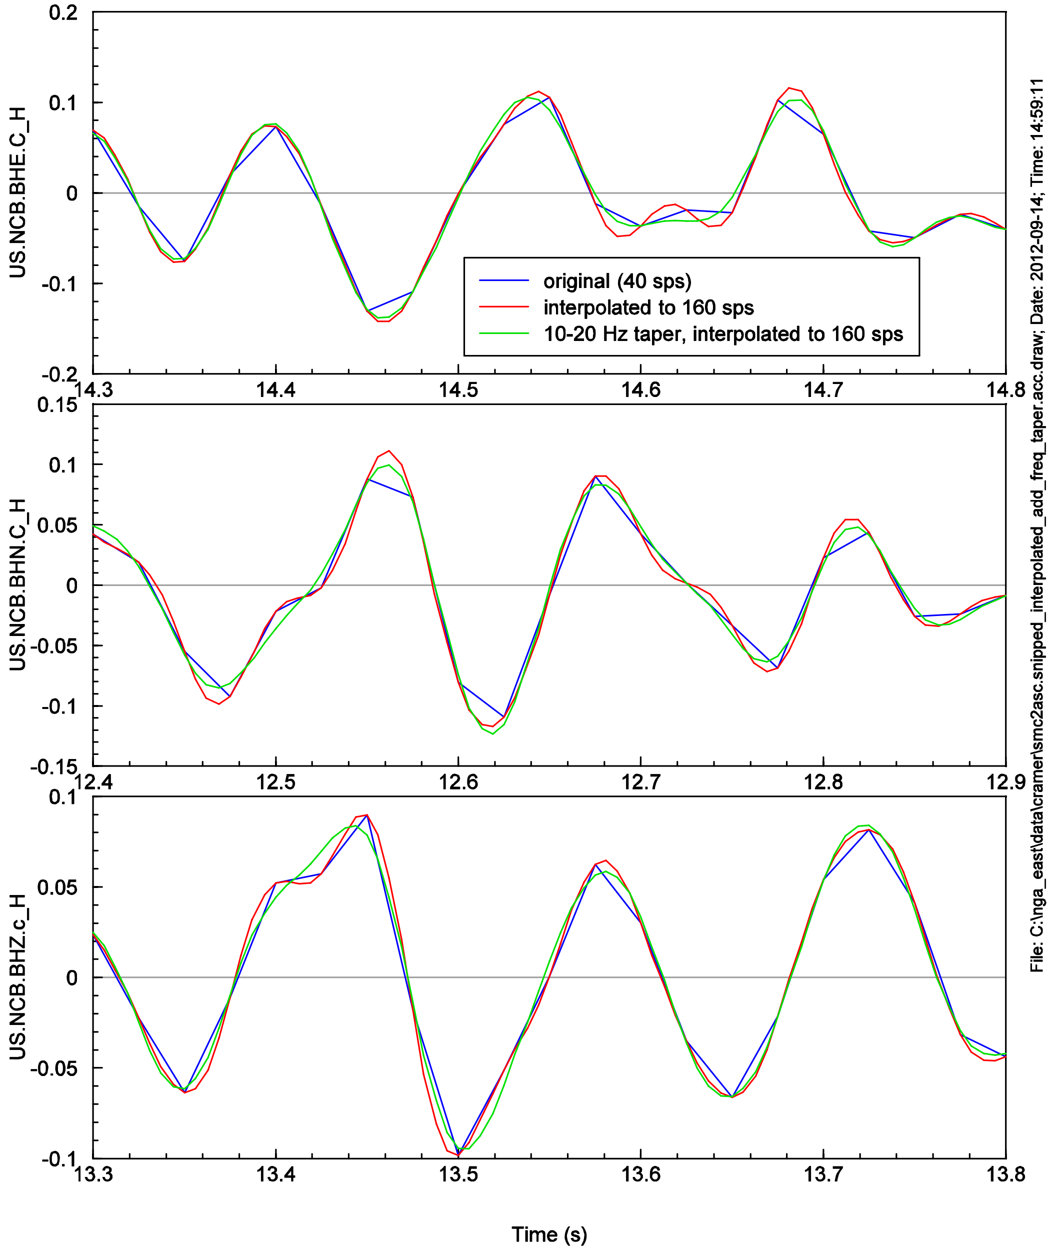

Figure 5. Original and resampled time series (the FAS taper from 10 to 20 s changes the time series, even without resampling, so that it no longer necessarily coincides with the original time series at the sample times.

Figure 6 shows the PSA for the three time series for each component, at enlarged scales compared to Figure 3.

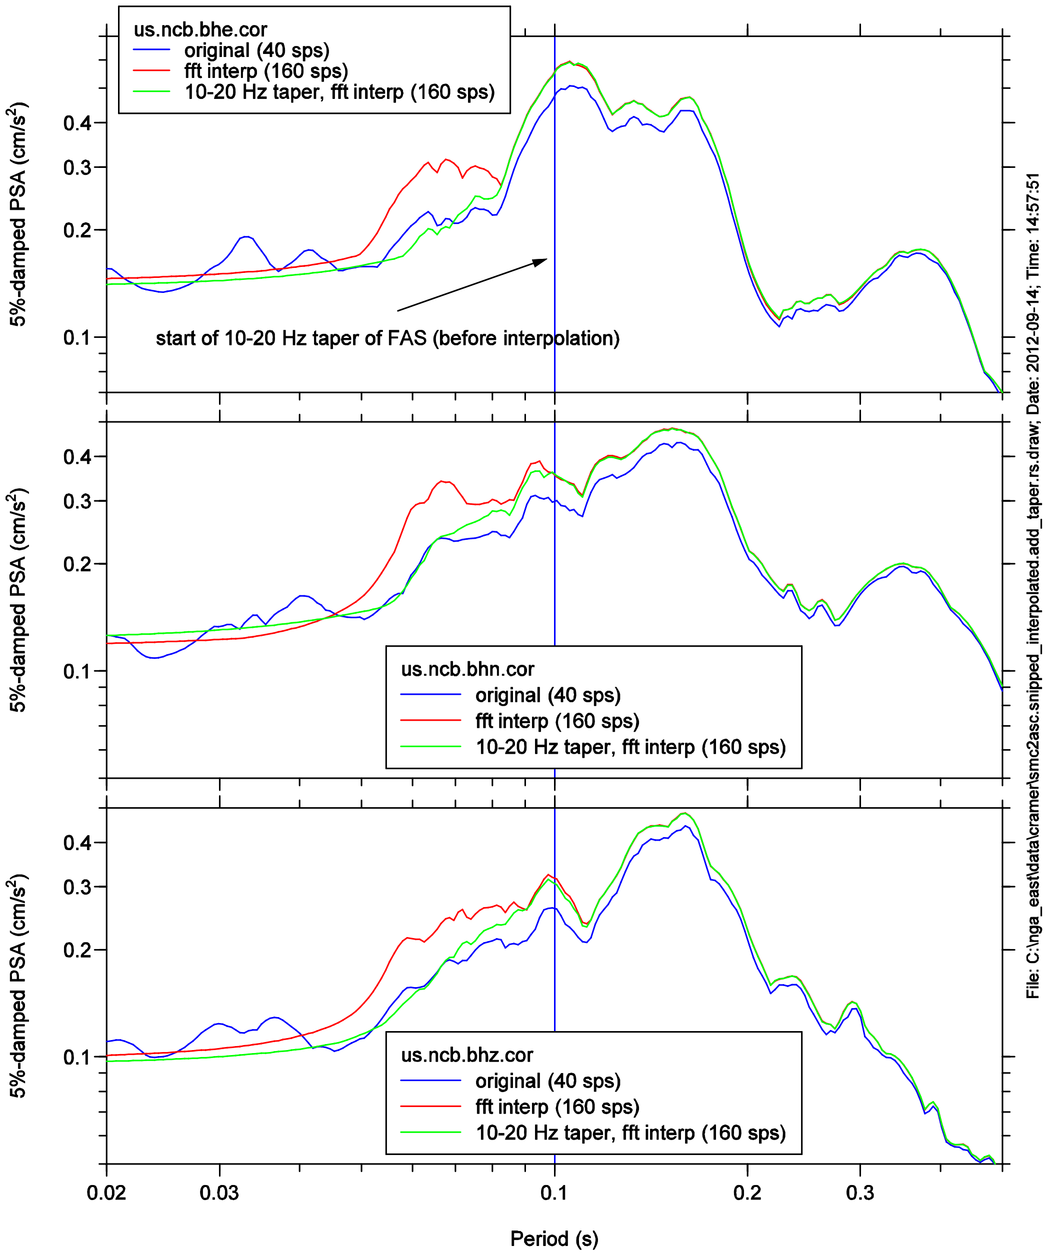

Figure 6. PSA for the three time series. The vertical blue line is corresponds to the lowest frequency of the 10 to 20 Hz taper applied to the FAS of the original records.

This figure shows that there is virtually no difference in the PSA from the resampled records for periods greater than the onset of the taper at 10 Hz (a period of 0.1 s). But the PSA from the original records are noticeably different than those from the resampled time series for periods as great as 0.3 s (the difference obviously decreases with increasing period). The differences at shorter periods are significant, and it is these short periods that are of particular importance in the response of nuclear power plants.

So I've shown that there can be a difference in the response spectra computed from time series with different assumptions about the behavior of the time series between the sampled points. But are the spectra from the resampled time series any closer to reality than from the original time series? One way of investigating this is by use of simulations, as was done by Douglas and Boore (2011) to study the effect of high-frequency noise on response spectra. I do not have time to do that now, but I include below a first look at the question, using data sampled at a higher rate, for which the FAS shows little decay at frequencies below the start of the anti-aliasing filter.

The data mentioned above were used in Douglas and Boore (2011; their Figure 3)). The sample rate is 125 *sps*; the FAS is shown in Figure 7.

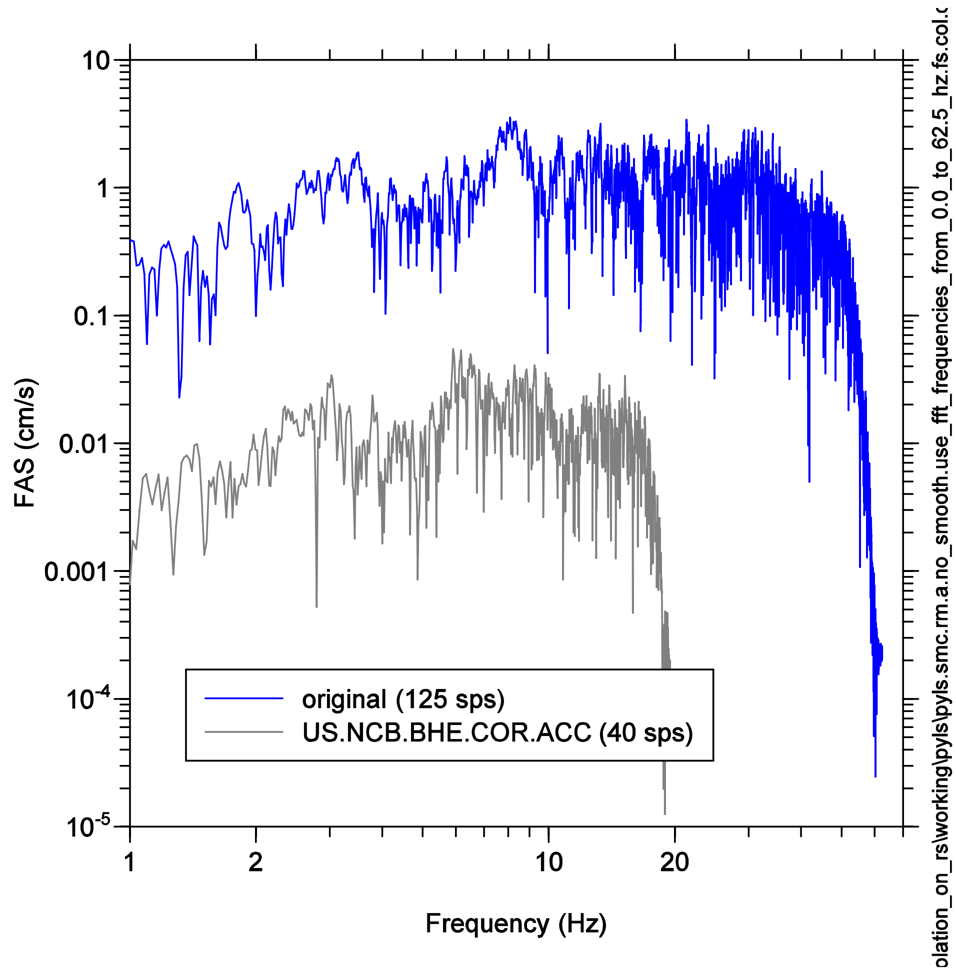

Figure 7. FAS of E components of data from PYLS (file name 2006.321.18.19.45.4349.RA.PYLS.00.ENE; see Figure 3 in Douglas and Boore, 2011)

For comparison, I also show in gray the FAS of the E component of the first example. As with the first example, the FAS shows little decay before the start of the anti-aliasing filter (a hardware filter, I think) at about 52 Hz, but the sample rate is about three times higher (125 *sps* compared with 40 *sps*). (The statement about the FAS being flat to the apparent start of the antialiasing filter is not quite true; a plot using a linear Y axis shows the FAS to be flat to about 30 Hz, after which it starts to decay at a rate of about f^{-2} to about 50 Hz, beyond which is decays much more rapidly). I resampled the time series as before (the new sample rate being 500 *sps*). A detailed portion of the original and resampled data is shown in Figure 8.

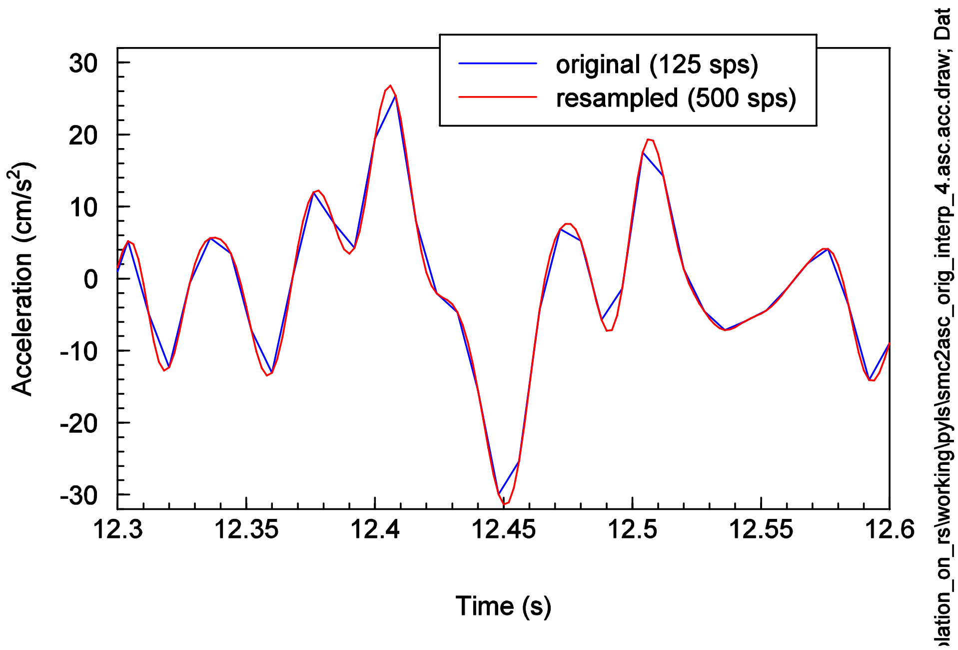

Figure 8. A small portion of the original and resampled acceleration in the vicinity of the absolute peak acceleration.

As before, the resampled time series smoothes the vertices and produces larger peaks. And as before, there are corresponding differences in the response spectra, as shown in Figure 9 (which includes the individual spectra as well as the ratio of the spectra).

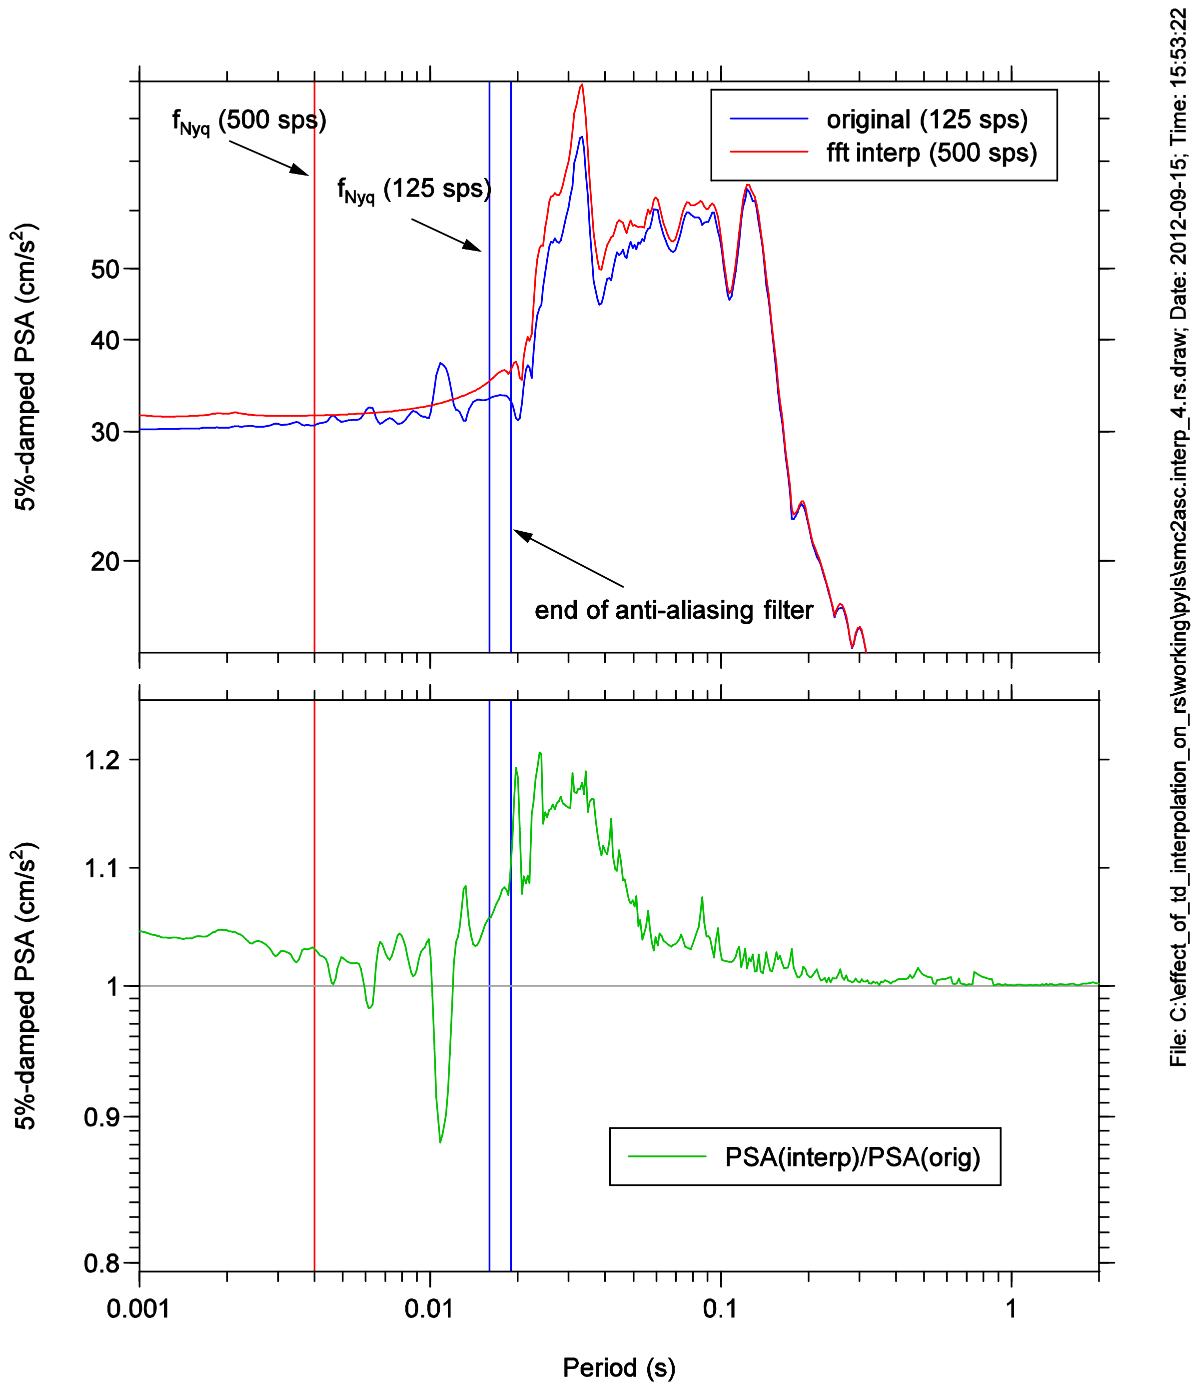

Figure 9. PSA and ratios of PSA for the PYLS days, without and with resampling.

As in the first example, there is a difference in PSA for periods greater than the start of the anti-aliasing filter, although the differences are not as large as before (a factor of 1.2 rather than 1.5—this may be due to the FAS starting to decay before the filter, whereas in the first example the FAS was more nearly constant to the start of the antialiasing filter). The two examples show that the influence of the time domain interpolation on response spectra is not fundamentally dependent on the Nyquist frequency, but rather on the shape of the FAS up to the Nyquist frequency.

The sampling theorem interpolation is obviously the theoretically proper way of interpolating between the sampled points, but this does not mean that the PSA from the resampled time series are equal to the real PSA for situations in which the FAS has little decay before the start of the anti-aliasing filter (this is similar to the conclusion of Douglas and Boore, 2011, regarding the relation between high frequency PSA and the shape of the FAS). As a first attempt to look at the question of whether the PSA from the resampled data are correct at periods close to the anti-aliasing period for cases when there is little decay in the FAS, I used the PYLS record. I applied a filter (a cosine taper) that approximated an anti-aliasing filter between 18 and 21 Hz; I then decimated the 125 *sps* record by a factor of 3 to obtain simulated the record that would have been obtained with a low sample rate (42 *sps*). I then computed the FAS and PSA for this record and this record resampled to a sample rate of 167 *sps*. Figure 10 compares the FAS of the original and high-cut filtered, decimated record.

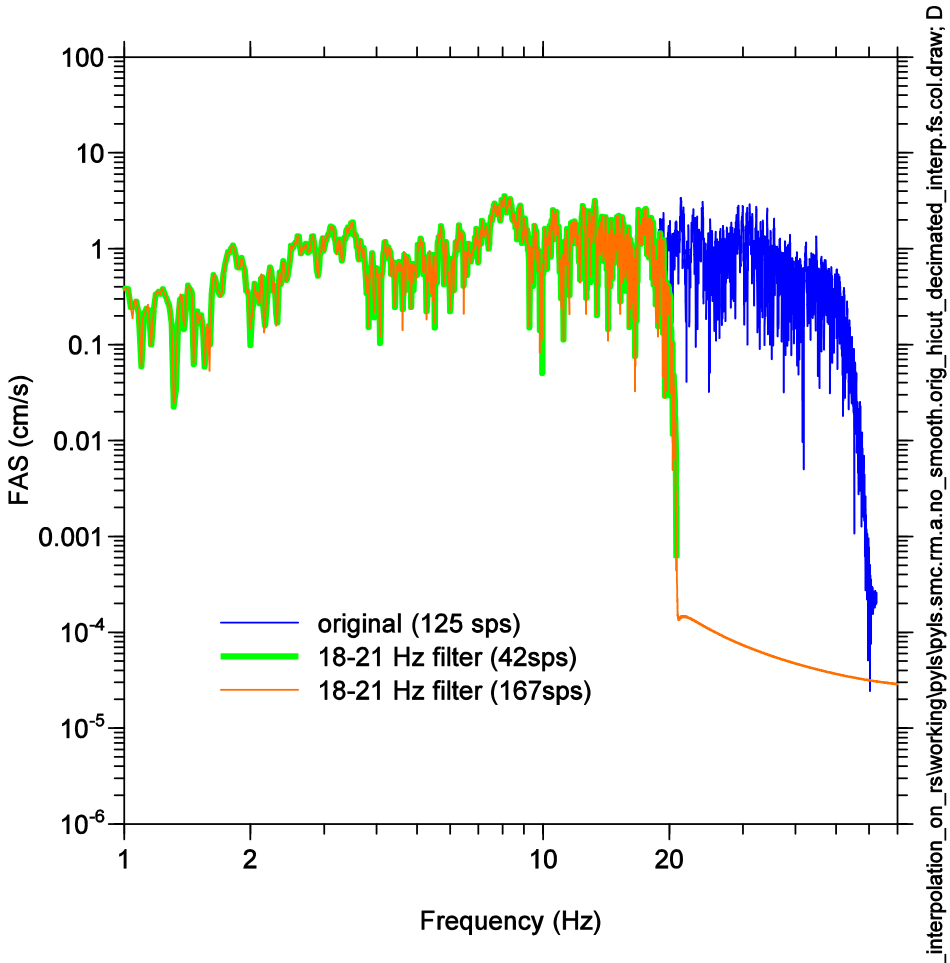

Figure 10. FAS of the original PYLS record and the simulated anti-alias filtered and low sample rate recorded record (and with the resampled version of that simulated record).

The PSAs and ratios of PSA are shown in Figure 11.

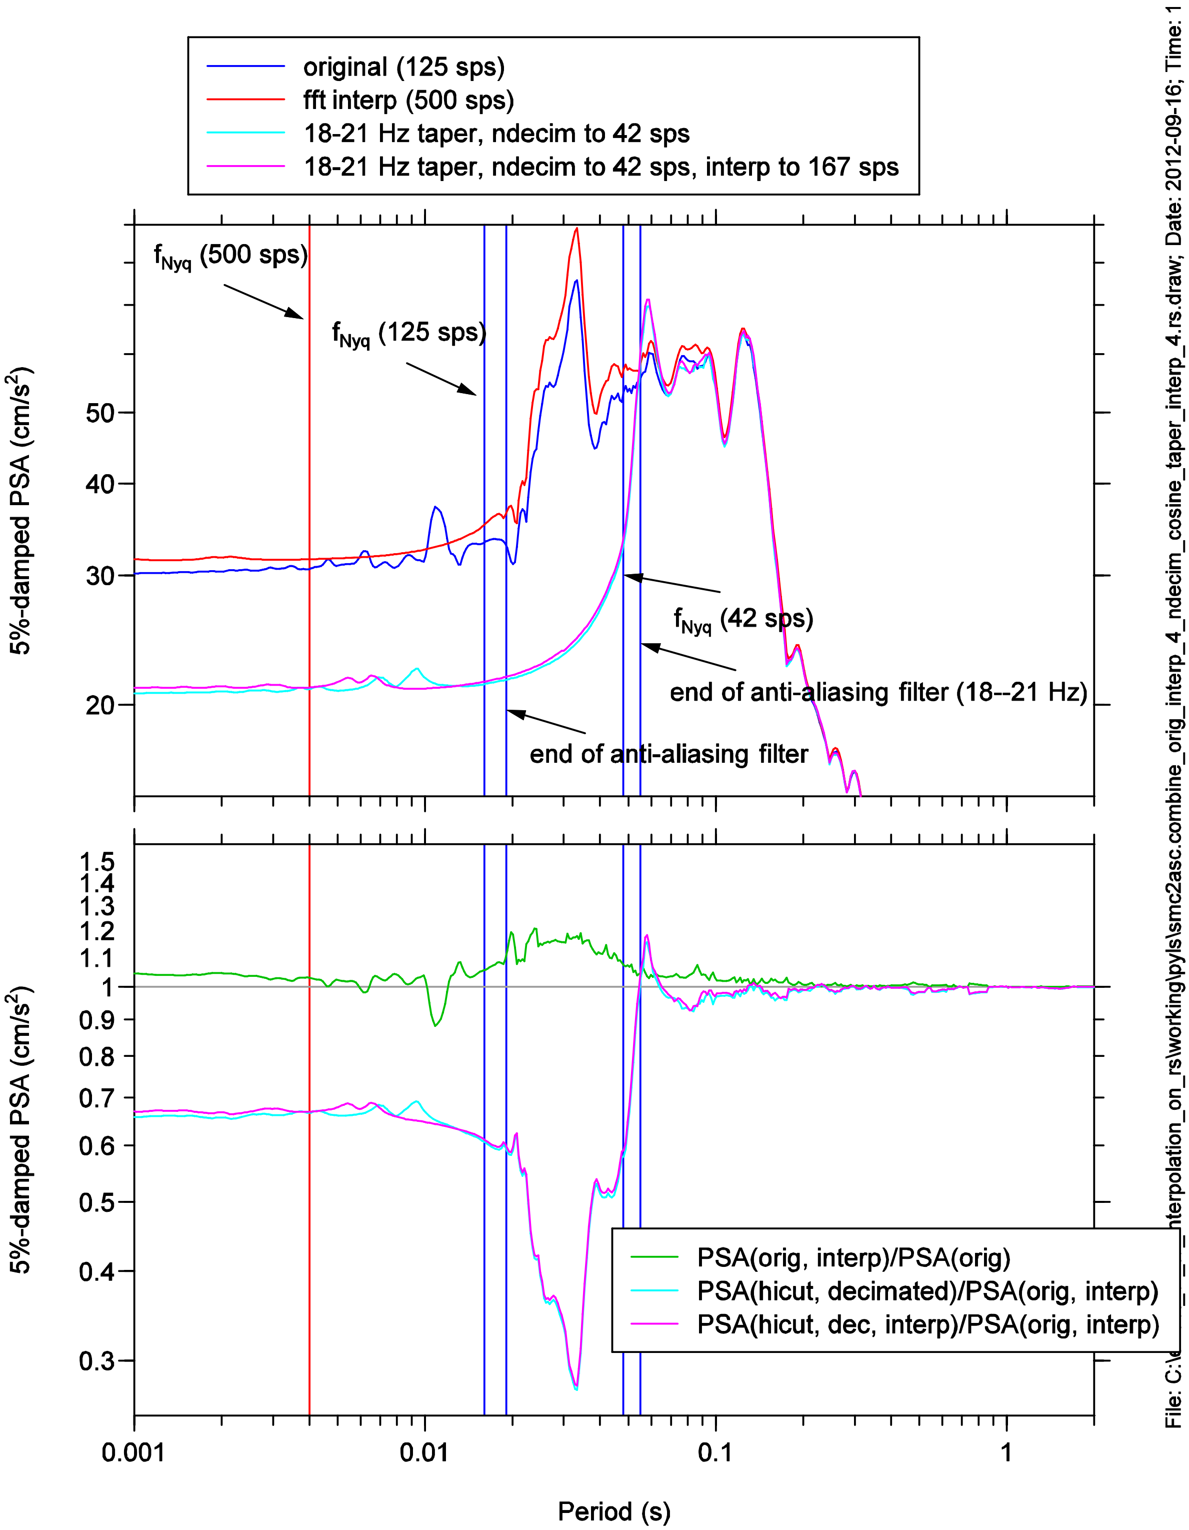

Figure 11. PSAs and ratios of PSA.

There is a lot of information here, but what I am interested in is the comparison of PSA from the simulation of anti-aliasing filtered, 42 *sps* data with the PSA from the best approximation of the original time series (with an anti-aliasing filter starting at about 52 Hz, resampled to 500 *sps*). I assume that the PSA from the original, resampled record can be taken as the true PSA for periods much lower than the original antialiasing filter (the transition bands of the anti-aliasing filters are shown in Figure 11). If so, then the ratio of PSA from the high-cut filtered and decimated record to the true PSA is not unity for periods close to the beginning of the anti-aliasing period. In particular, there is a spike in the ratio with an amplitude close to 1.2. It may be that this is a result of applying a sharp anti-aliasing filter (i.e., one with a narrow transition band) to ground motion whose FAS is nearly flat (see Figure 10). It would be interesting to look into this in more detail using simulations, but for now the results suggest that PSA for periods near the start of the anti-aliasing period could be biased, whether or not the resampling to a higher sample rate is performed. The crucial factor of whether or not this is important is the behavior of the true FAS near the anti-aliasing transition band. If the FAS has decayed appreciably for frequencies approaching the anti-aliasing frequency (presumably because of natural attenuation in the earth), then the PSA will probably be a good approximation of the true PSA, even for very short periods (e.g., Douglas and Boore, 2011), and no resampling of the time series will be necessary. For situations when this is not true, the PSA from the original and resampled data will differ, but it is not clear whether the PSA from the resampled data will be a good apprioximation of the true PSA for periods near the anti-aliasing transition band. More study of this is needed.

References

Douglas, J. and D. M. Boore (2011). High-frequency filtering of strong-motion records, *Bull. Earthquake Engineering* **9**, 395—409.

Press, W.H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery (1992). *Numerical Recipes in FORTRAN: The Art of Scientific Computing*, Cambridge University Press, Cambridge, England, 963 pp.