


Gibbs Oscillations in Bandlimited Frequency Response Functions

Brad Thoen, Staff Engineer
MTS Systems Corporation
Rev. 18-Aug-14

A common technique of filtering in the frequency domain is to zero the frequency bins associated with unwanted content. Unfortunately the discontinuities in the frequency response function (FRF) caused by the zeroing operation give rise to significant oscillations between frequency bin samples. These oscillations, known as Gibbs oscillations, can be mitigated by smoothing the FRF with a low order window so that the passband-stopband transitions are not so abrupt.


Gibbs oscillations were discovered by a 19th century mathematician of that name when he tried to approximate a discontinuous function with a finite number of terms of a Fourier series. A classic exercise known to all engineering undergraduates is approximating a square wave with a Fourier series:




The more terms that are included, the closer the approximation approaches a square wave. Note that the peaks do not decrease in amplitude with the number of terms, although they do get closer to the discontinuity and the area beneath them decreases. This is the phenomenon of Gibbs oscillations, and results from approximating a discontinuity with a basis of continuous sinusoids.

The same effect occurs when bandlimiting an FRF by zeroing unwanted frequency bins. It will now be demonstrated that the zeroing process creates discontinuities in the FRF, giving rise to Gibbs oscillations. It will then be shown how to mitigate these oscillations.


In the example below, the system FRF will be bandlimited by applying a *boxcar* window, which preserves the lower frequency bins and zeros the upper frequency bins:



After multiplying the system FRF by the boxcar FRF, there is now a significant discontinuity in the resulting bandlimited system FRF:




It is important to realize that the FRF represented by the dots are values of a continuous FRF that has been sampled on a frequency grid, i.e., *there is frequency content between frequency samples*. In order to see this content, upsample the FRF to create a denser frequency grid, and the Gibbs oscillations become apparent:




If this FRF were only used to filter signals that are periodic and of the same sample rate, then what happens between frequency samples would be irrelevant, because the only frequencies that exist in such a signal are those that are on the frequency grid. This is not true of aperiodic signals, which have content at a continuum of frequencies. Most signals of interest in the real world are aperiodic, so what happens between frequency bins does matter. Also, if the filtered signal is sent out to the real world via digital-to-analog conversion, the signal is upsampled to continuous time, and once again inter-bin frequency content matters.

So for real world signals, the filter represented by the FRF above is a poor one: there is significant content in the stopband, and the passband has significant ripple. But it is possible to improve this FRF for filtering purposes by reducing the Gibbs oscillations, using the method described next.

Gibbs oscillations between frequency bins can be reduced by smoothing the abruptness of the discontinuity. This is done by applying a low-order averaging window to the FRF in the frequency domain, which creates additional points in the transition between passband and stopband, where before there were none:



Upsampling the smoothed FRF, we find that the Gibbs oscillations are considerably reduced:



The smoothing window can be applied multiple times to get an even smoother transition and consequently smaller Gibbs oscillations. However, because this smoothing may also change the passband in undesirable ways, it is best not to overdo it.