CS 583– Computational Audio -- Fall, 2021

Wayne Snyder
Computer Science Department
Boston University

Lecture 10

The Discrete Fourier Transform (DFT) and the Fast Fourier Transform (FFT)

Issues/Problems with the DFT and (partial) solutions

Computer Science

Discrete Fourier Transform (DFT)

Define: For a signal X of length W samples (a "window") a window frequency is one whose period P is such that W = P * k for some integer k, i.e., an integral number of periods exactly fit within the window; alternately, it begins and ends at same instantaneous phase.

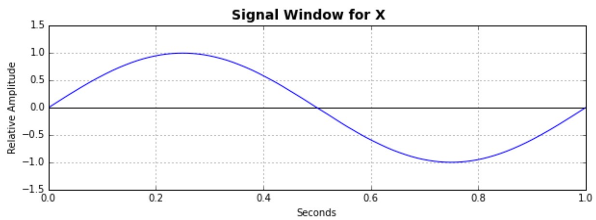

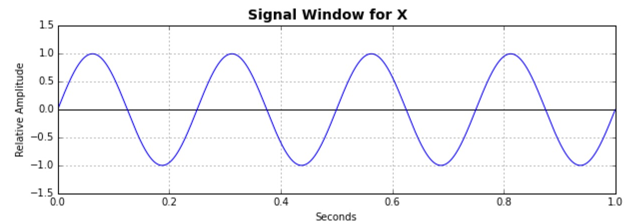

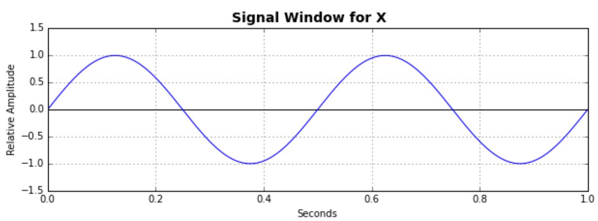

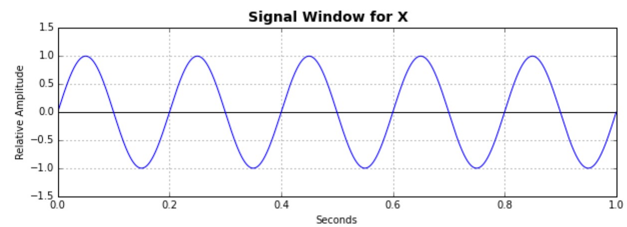

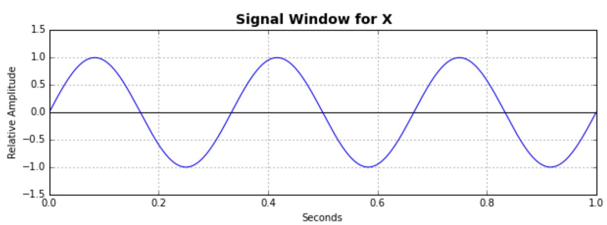

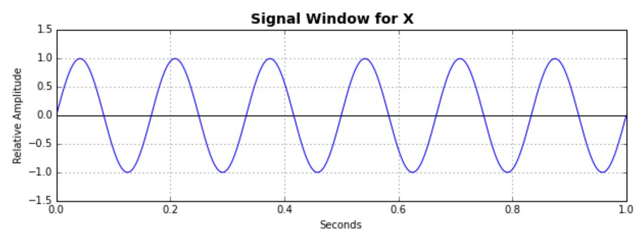

We will use these signals as **probe waves** to analyze a musical signal and assume that all such probe waves (for now) start at phase 0.0.

Using the DFT

The way we will analyze a signal using the DFT is to examine a short segment of a signal (a window):

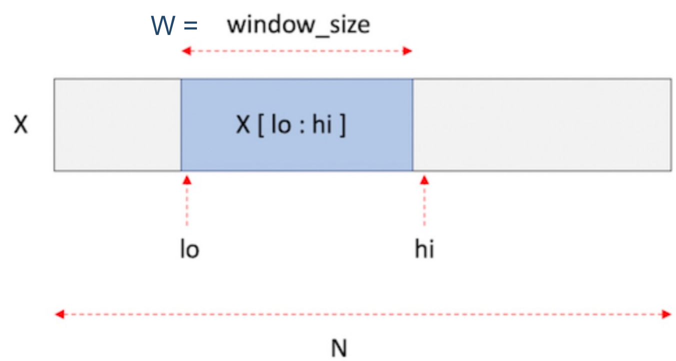

The FFT uses a recursive "divide and conquer" strategy, and so it is best if W is a power of 2.

If we wish to examine the whole signal, we will slide the window across the signal, potentially overlapping each window:

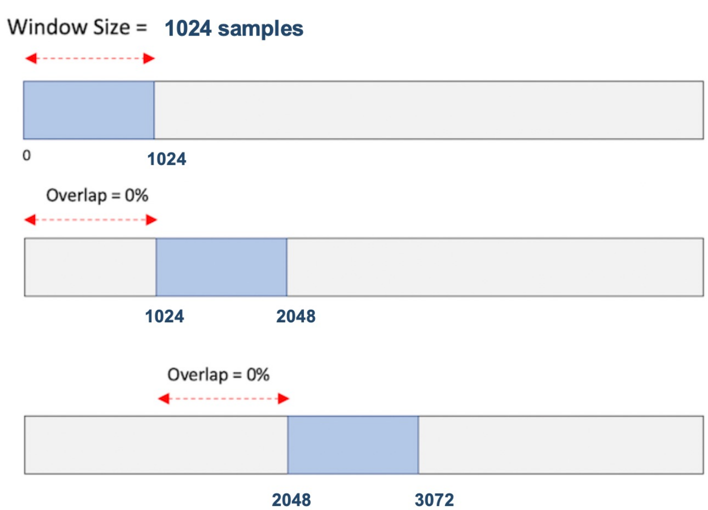

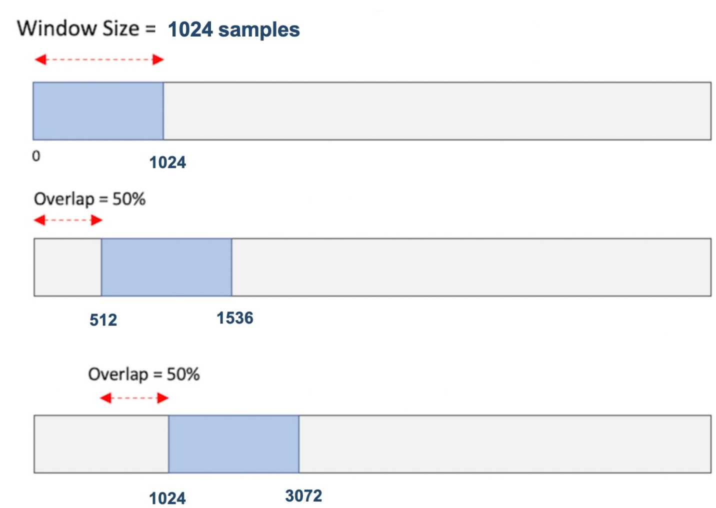

Using the DFT

Computer Science

The DFT produces a spectrum for a window (erroneously called the "instantaneous spectrum"), containing the amplitudes of frequencies

Bin index: 0 1 2 3 $\lfloor W/2+1 \rfloor$
Frequency: 0, f, 2f, 3f, …., Sample Rate / 2 = Nyquist Limit

f = frequency of one cycle per window = Sample Rate / W
= frequency resolution

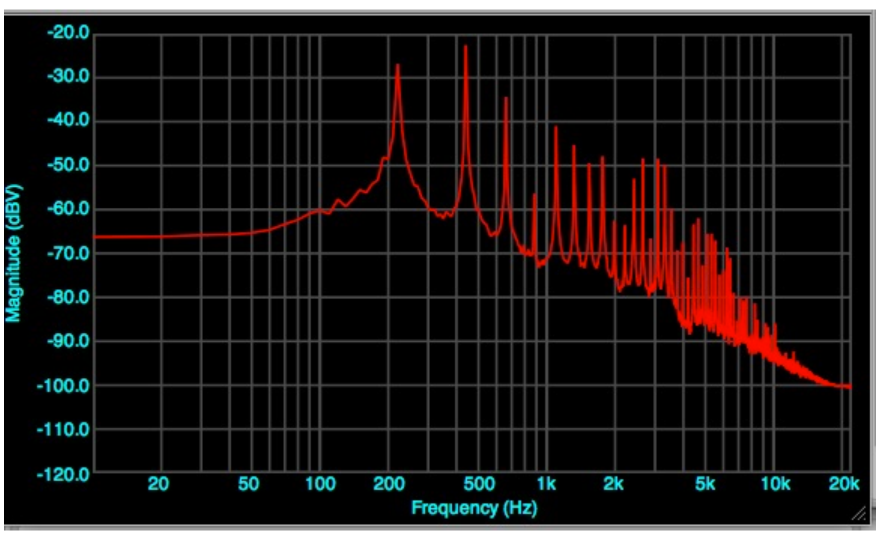

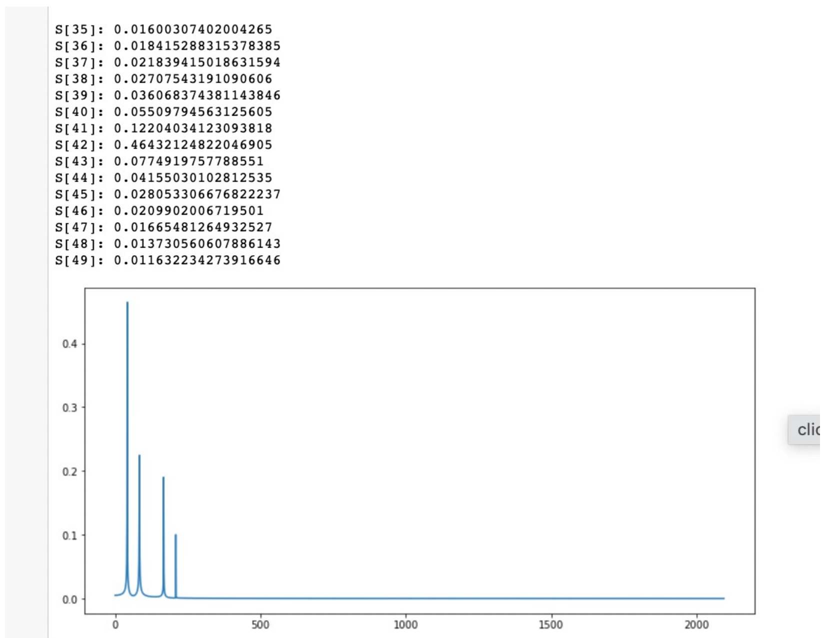

Computer Science

The axes for the instantaneous spectrum are

X axis = Frequency

Y axis = Amplitude ("magnitude spectrum")

or

Amplitude ^{2} ("power spectrum")

Either one could be displayed in a linear or a logarithmic scale; the logarithmic scales more closely represent the way humans perceive both frequency and loudness. Decibels (dB) are a logarithmic measure.

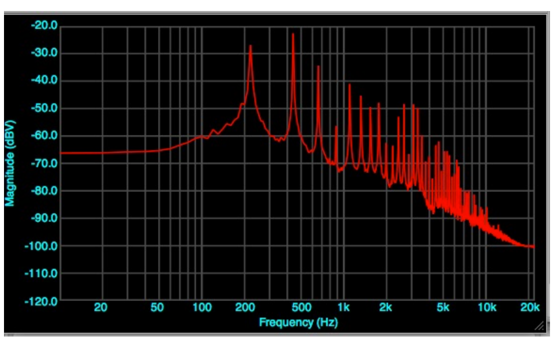

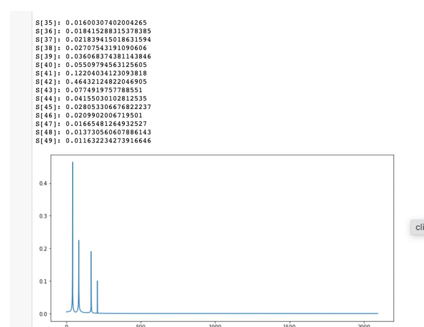

The FFT uses linear scales for both!

Using the DFT

Computer Science

A spectrogram is a 2D matrix, usually presented as a colored "heatmap" with spectra in each column for each sliding window (possibly overlapping), and rows being the frequency bins.   The axes are thus

X axis = Time   (always linear, usually in seconds)

Y axis = frequency   (either linear or logarithmic (dBs))

Color = amplitude   (magnitude (linear) or power (squared amplitude))

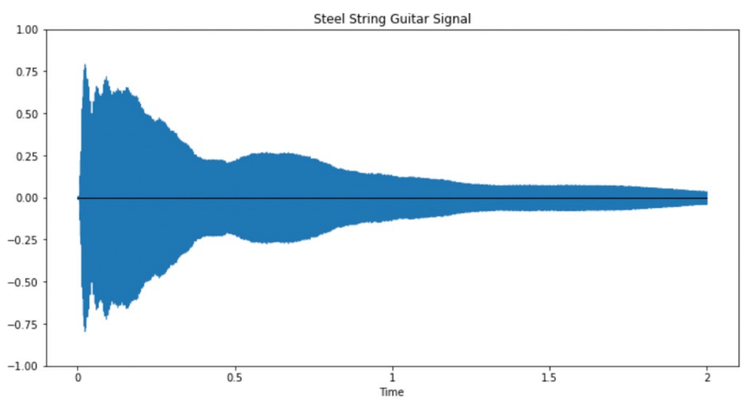

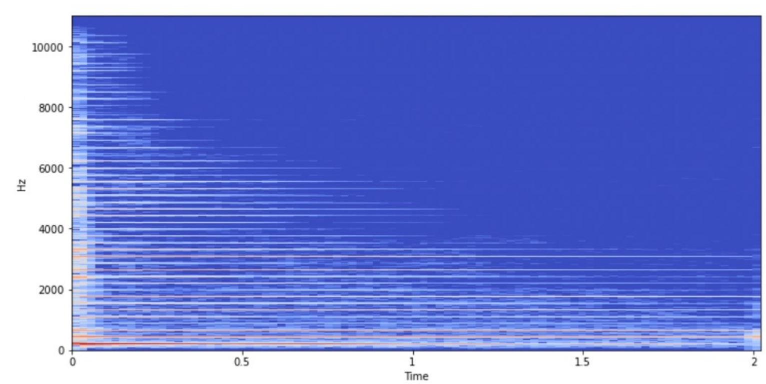

Using the DFT

Computer Science

Sometimes spectrograms are presented in faux-3D with or without heatmap colors:

X axis = Time (always linear, usually in seconds)

Y axis = frequency (either linear or logarithmic (dBs))

Z axis = amplitude (magnitude (linear) or power (squared amplitude))

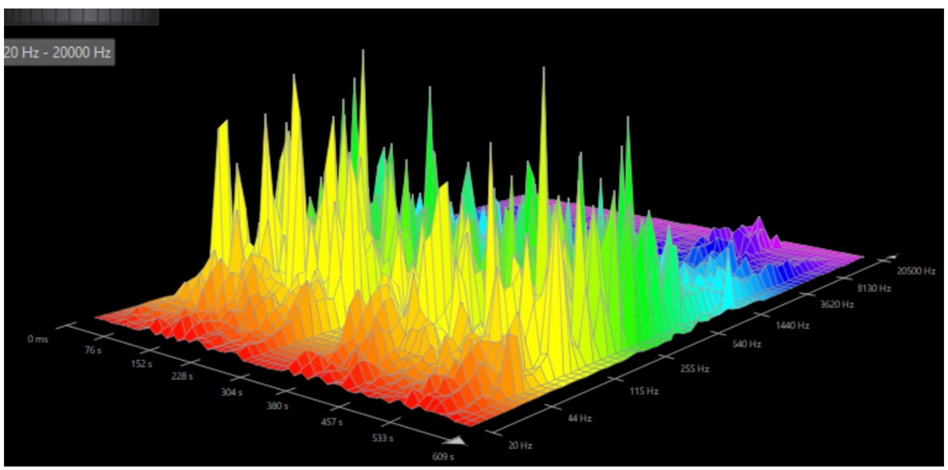

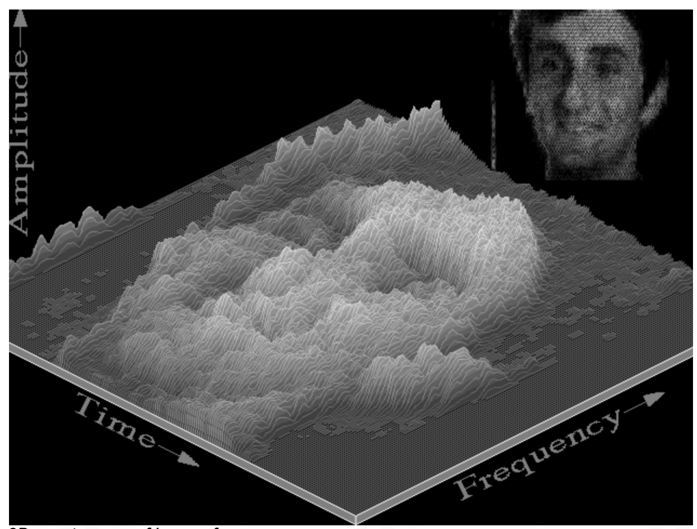

3D spectrogram of human face.

Discrete Fourier Transform (DFT)

Recall from last time:

(1) This is horribly inefficient: O(N^2) for N = len(X)

✓ Solution: There is a fast version of the transform, the Fast Fourier Transform (FFT), based on a recursive algorithm, which runs in O(N log(N)).

(2) The resolution is limited to multiples of f = SR / W (in samples)

✗ No solution, unfortunately, can try different window sizes, but stuck with this!

(3) All components and probe waves **have to be at the same phase (e.g., 0.0)**

✓ Solution: If we do all the work with **complex numbers**, we can avoid issues of phase

Limitations and Problems with the DFT

Unfortunately, this is not the only limitation of the DFT. Here are the main issues we need to be aware of when using the DFT:

1. The resolution is limited to multiples of f = SR / W (in samples), and so there is a tradeoff (the “DFT Uncertainly Principle”) between temporal resolution and frequency resolution.

2. When frequencies in the signal do not exactly correspond to the window frequencies, their energy is spread out among the closest frequency bins, so the amplitude is not represented precisely.

DFT: Tradeoffs in Resolution of Frequency and Time

Computer Science

There is a tradeoff between

Temporal Resolution – What is the shortest musical event we can observe?
Spectral Resolution – How many frequencies can we measure?

← Window of W Samples →

There is a tradeoff between

Temporal Resolution – What is the shortest musical event we can observe?
Spectral Resolution – How many frequencies can we measure?

← Window of W Samples →

The duration of the window is W / SR, e.g., if SR = 22050, then

W	Time Resolution
64	0.0029
128	0.0058
256	0.0116
512	0.0232
1024	0.0464
2048	0.0929
4096	0.1858
8192	0.3715

Significance of temporal resolution for musical signals

30 second sample of Bob Marley…..

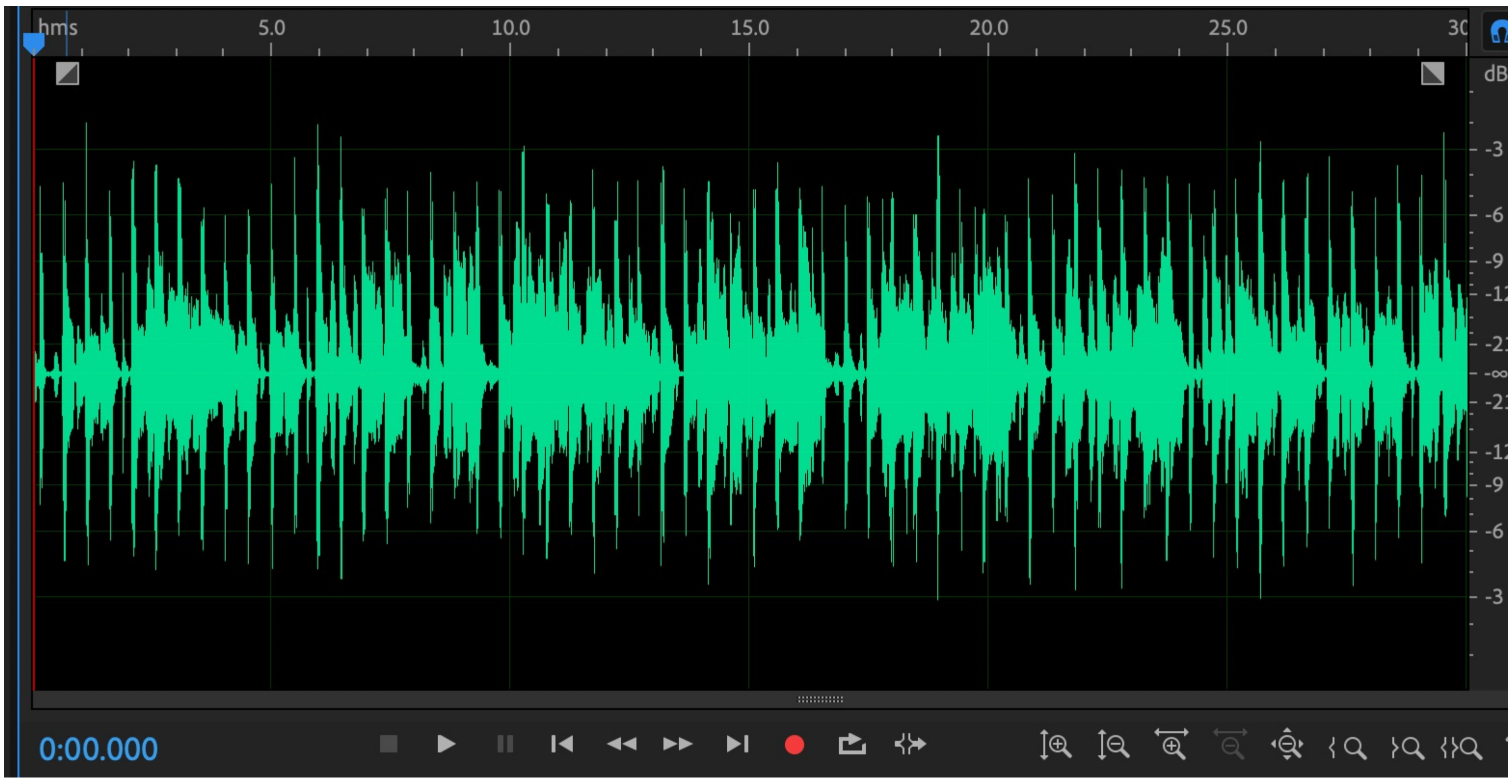

Significance of temporal resolution for musical signals

1 second window:

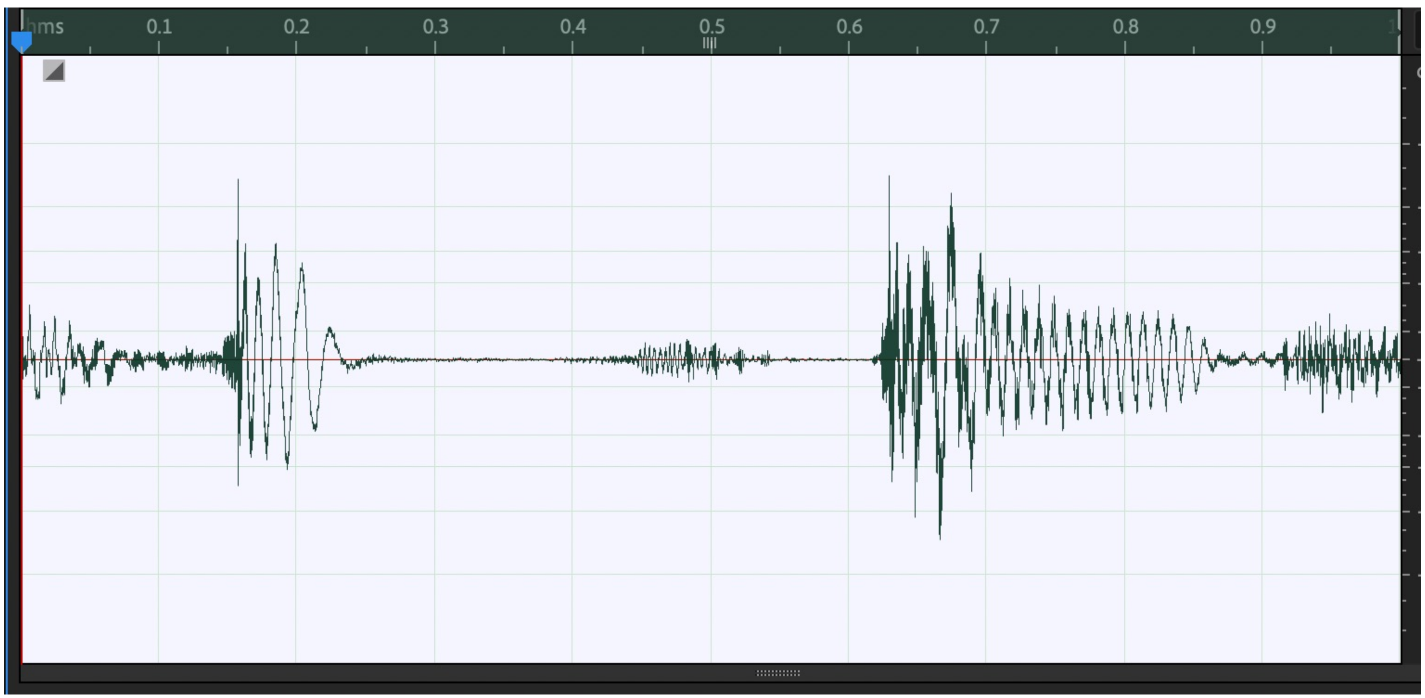

Significance of temporal resolution for musical signals

0.2 second window:

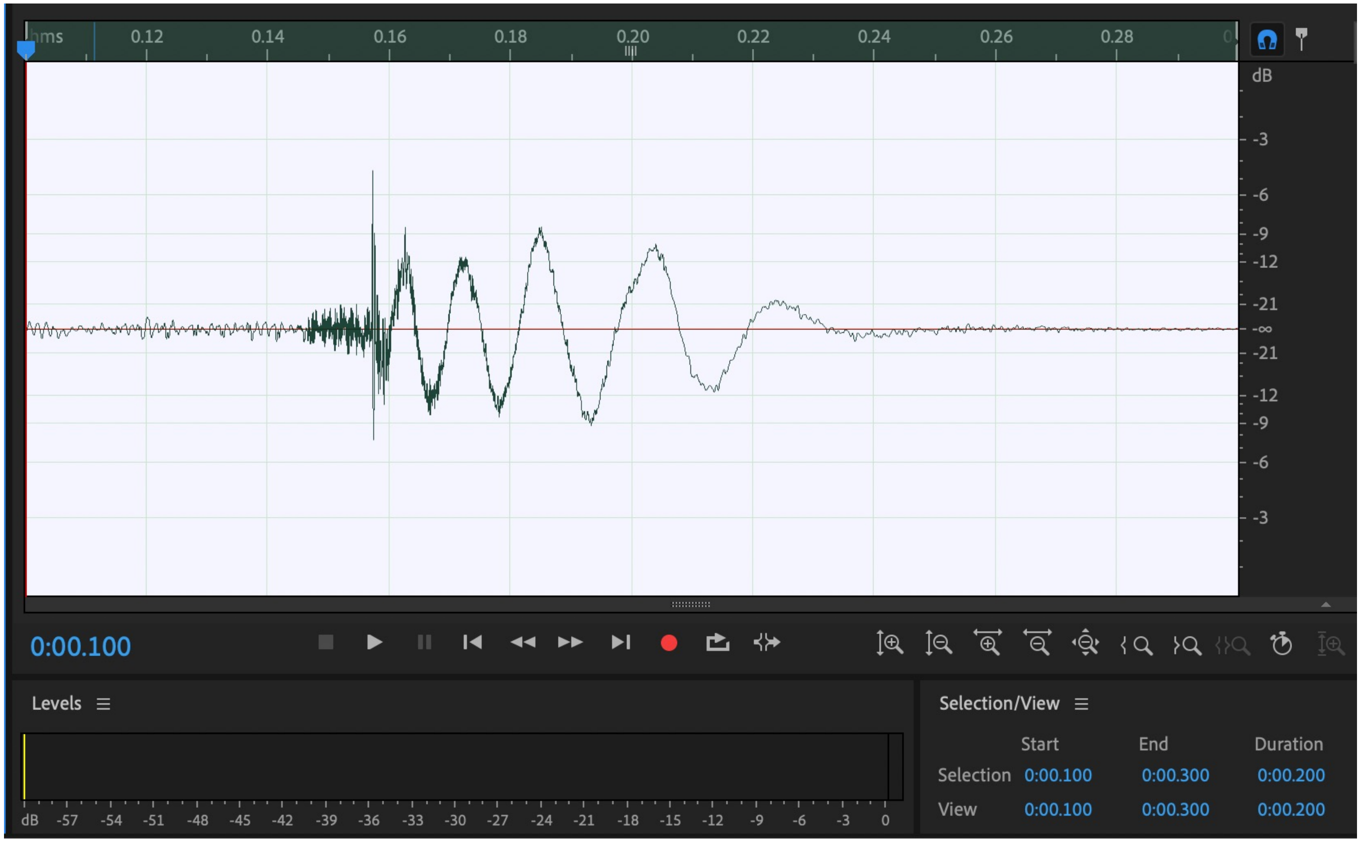

Significance of temporal resolution for musical signals

0.02 second window:

Significance of temporal resolution for musical signals

0.002 second window:

Recall: There is a tradeoff between

Temporal Resolution – What is the shortest musical event we can observe?

Spectral Resolution – How many frequencies can we measure?

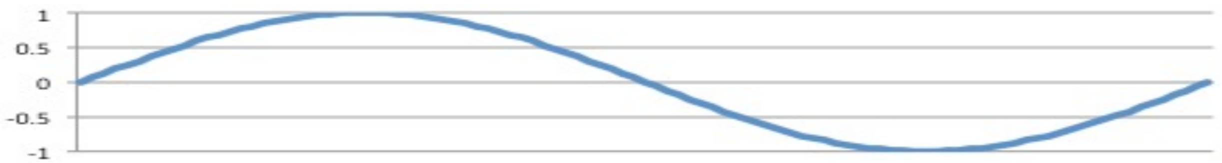

← Window of W Samples →

But then temporal and frequency resolution are in an inverse relationship:

W	Time Resolution	Frequency Resolution
64	0.0029	344.5312
128	0.0058	172.2656
256	0.0116	86.1328
512	0.0232	43.0664
1024	0.0464	21.5332
2048	0.0929	10.7666
4096	0.1858	5.3833
8192	0.3715	2.6917

Frequency vs temporal resolution

W	Time Resolution	Frequency Resolution
64	0.0029	344.5312
128	0.0058	172.2656
256	0.0116	86.1328
512	0.0232	43.0664
1024	0.0464	21.5332
2048	0.0929	10.7666
→ 4096	0.1858	5.3833
8192	0.3715	2.6917

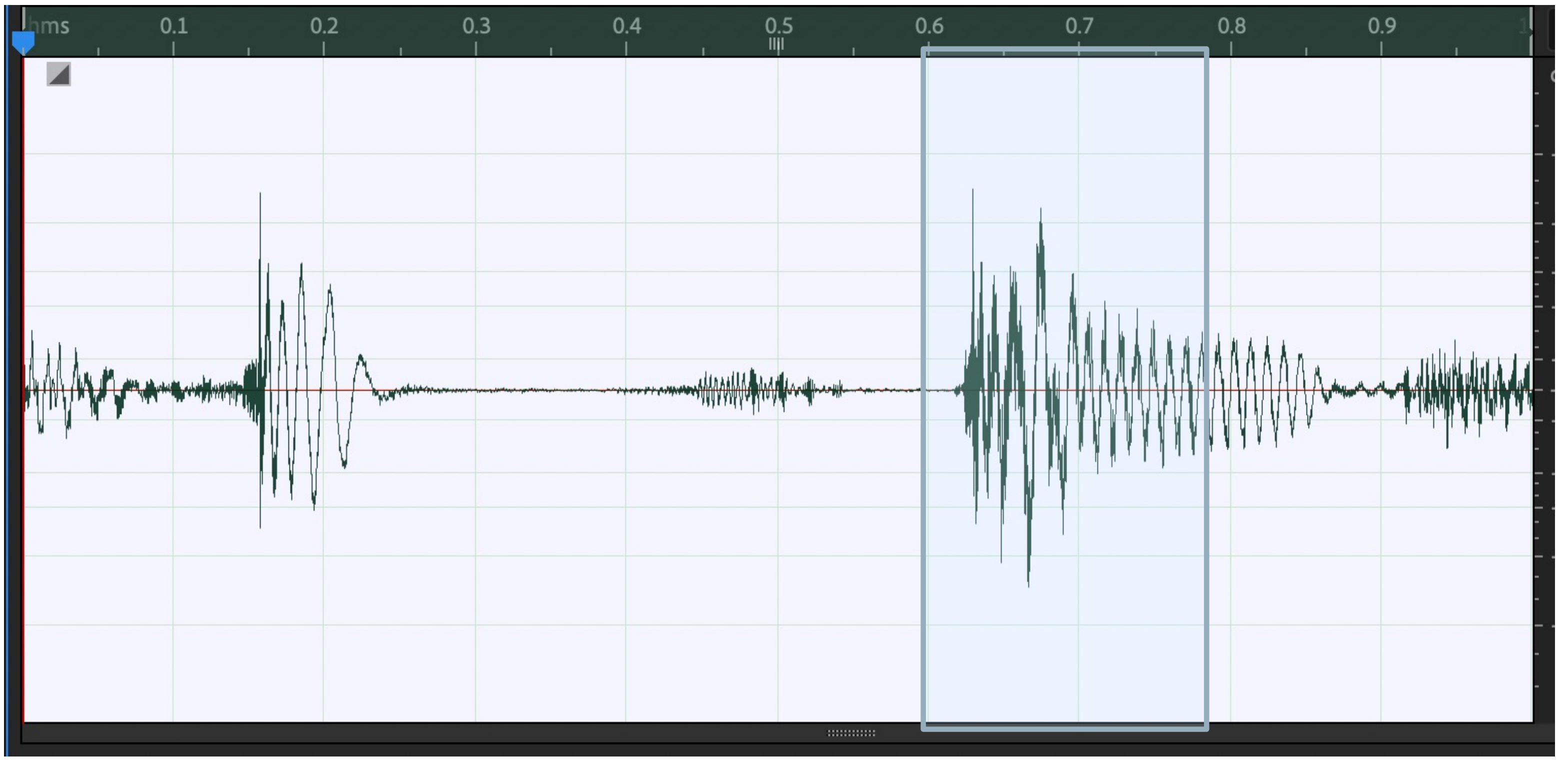

Digital Audio Fundamentals: The Discrete Fourier Transform

Frequency vs temporal resolution

W	Time Resolution	Frequency Resolution
64	0.0029	344.5312
128	0.0058	172.2656
256	0.0116	86.1328
512	0.0232	43.0664
1024	0.0464	21.5332
2048	0.0929	10.7666
4096	0.1858	5.3833
8192	0.3715	2.6917

Digital Audio Fundamentals: The Discrete Fourier Transform

Frequency vs temporal resolution

W	Time Resolution	Frequency Resolution
64	0.0029	344.5312
128	0.0058	172.2656
256	0.0116	86.1328
512	0.0232	43.0664
1024	0.0464	21.5332
2048	0.0929	10.7666
4096	0.1858	5.3833
8192	0.3715	2.6917

Digital Audio Fundamentals: The Discrete Fourier Transform

Frequency vs temporal resolution

W	Time Resolution	Frequency Resolution
64	0.0029	344.5312
128	0.0058	172.2656
256	0.0116	86.1328
512	0.0232	43.0664
1024	0.0464	21.5332
2048	0.0929	10.7666
4096	0.1858	5.3833
8192	0.3715	2.6917

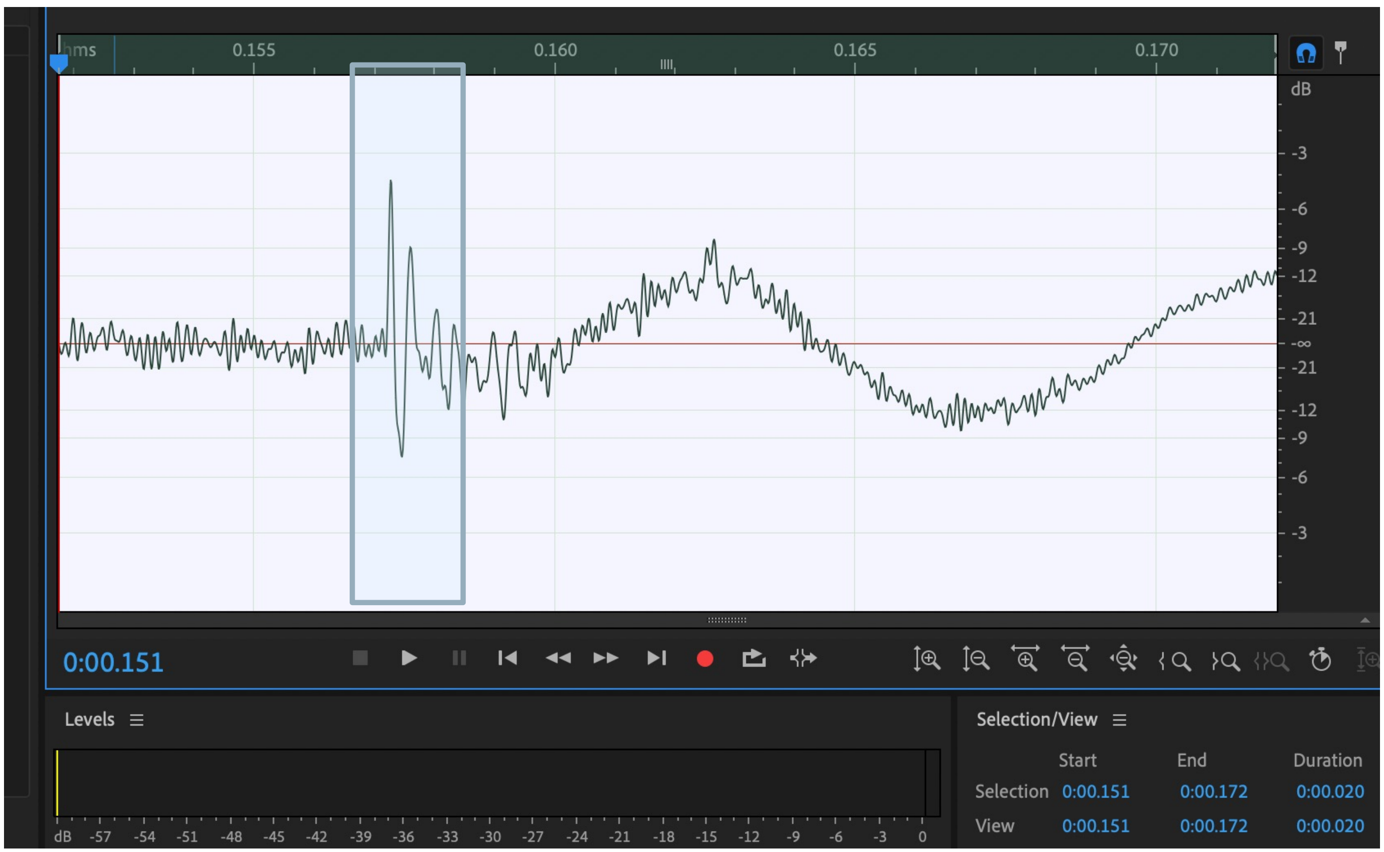

Assume that the sample rate is 44100 (CD quality sound).

To cover the range of human hearing (20 – 20,000 Hz), then, we would need a window of 44100/20 = 2205 samples, and this would give us the ability to measure frequencies

20, 40, 60, …., 22000 (Note that the upper bound is always the Nyquist Limit)

and the time resolution of this window is 0.05 = 1/20 sec.

To measure down to C2 (65.41 Hz, two octaves below Middle C) we would need a window of 44100/65.41 = 674.2 samples, with a resolution of 0.015 = 1/65.41 sec.

To measure down to E2 (82.41 Hz, the low string on a guitar) we would need a window of 44100/82.41 = 535.13 samples, with a resolution of 0.00186 = 1/82.41 sec.

Punchline: Probably for reasonable musical signals we have enough temporal resolution and the RANGE of the frequencies seems enough….

But then there is a problem with frequency resolution:

To cover the range of human hearing (20 – 20,000 Hz), then, we would need a window of 44100/20 = 2205 samples, and this would give us the ability to measure frequencies

20, 40, 60, 80, 100, 120, 140, 160, 180, 200, 220, …., 22050

For the C-major scale one octave above middle C we have the following:

523.25 554.37 587.33 622.25 659.26 698.46

520 540 560 580 600 620 640 660 680

PunchLine: We can't even come close to measuring all frequencies in a musical signal with one window size: we don't have enough frequencies and they don't match precisely the pitches….

So we could use different window sizes.....

Note	Freq	Window Size	Duration	Window Freq	Freq Detected	Error
C	32.7032	2697	0.061	2	32.7030	0.000197
C#	34.6478	5091	0.115	4	34.6494	0.001581
D	36.7081	3604	0.082	3	36.7092	0.001112
Eb	38.8909	2268	0.051	2	38.8889	0.002011
E	41.2034	3211	0.073	3	41.2021	0.001282
F	43.6535	4041	0.092	4	43.6526	0.000939
F#	46.2493	1907	0.043	2	46.2507	0.001355
G	48.9994	900	0.020	1	49.0000	0.000600
Ab	51.9131	1699	0.039	2	51.9129	0.000210
A	55	4811	0.109	6	54.9990	0.001039
Bb	58.2705	3784	0.086	5	58.2717	0.001170
B	61.7354	2143	0.049	3	61.7359	0.000484
A	220	2205	0.050	11	220.0000	0.000000

BUT, with one window, we can measure the multiples of a particular fundamental…. And this matches the way musical instruments work (in general):

A guitar tone of pitch 220 Hz (A below middle C, A on the G string) has its strongest components at the harmonics:

220, 440, 660, 880, ……,

And we can find these (luckily enough!) with a window size of 2205:

44100 * 11 / 2205 = 220

So 220 Hz ought to show up as frequency 11,
440 Hz as 22, etc.

Let's look.....

But it is not clear that the only frequencies are multiples of the fundamental, and each “peak” is not a simple value, but a ”triangular mountain”:

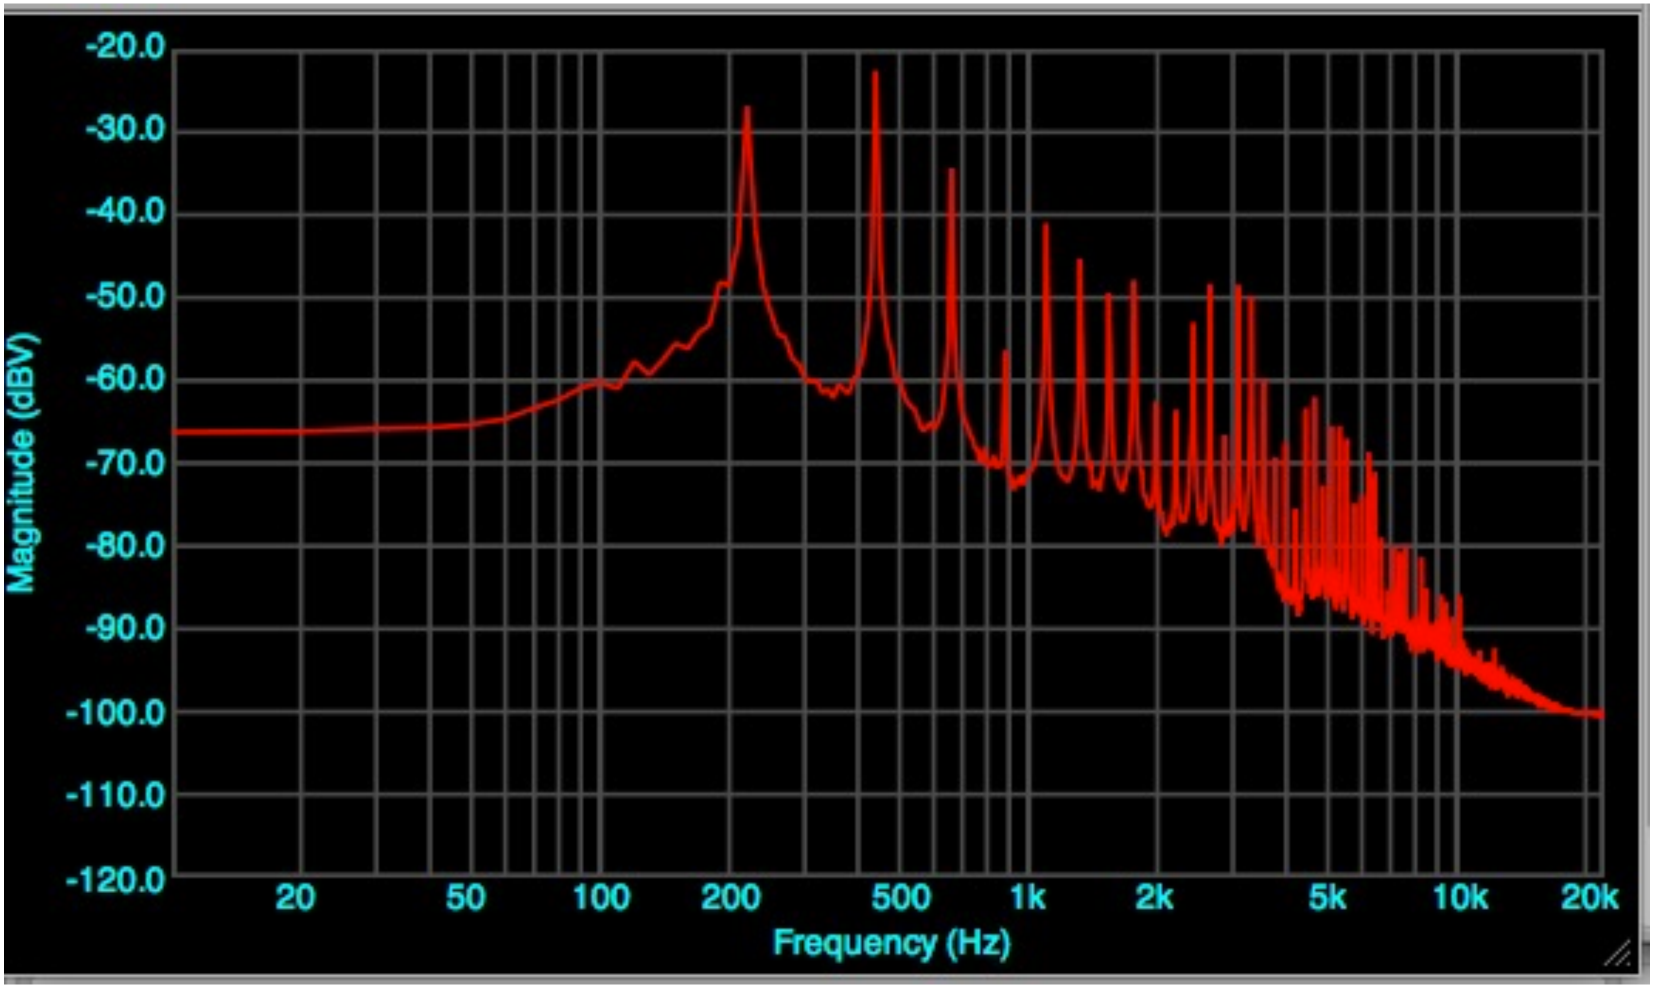

Let's explore why: When frequencies are integral (i.e., K complete periods within the window of N samples), we get precise measurements. Let's consider what happens with frequencies around 50 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

Exact Measurements around 50 Hz

	Component Waves								Signal
	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave
Freq:	50	4	6	12	20	56.14	78	99	
Amp:	1	0	0	0	0	0	0	0	1
Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3	

Freq:

Probe Freq	Amplitude Measure
47	0.0000
48	0.0000
49	0.0000
50	1.0000
51	0.0000
52	0.0000
53	0.0000

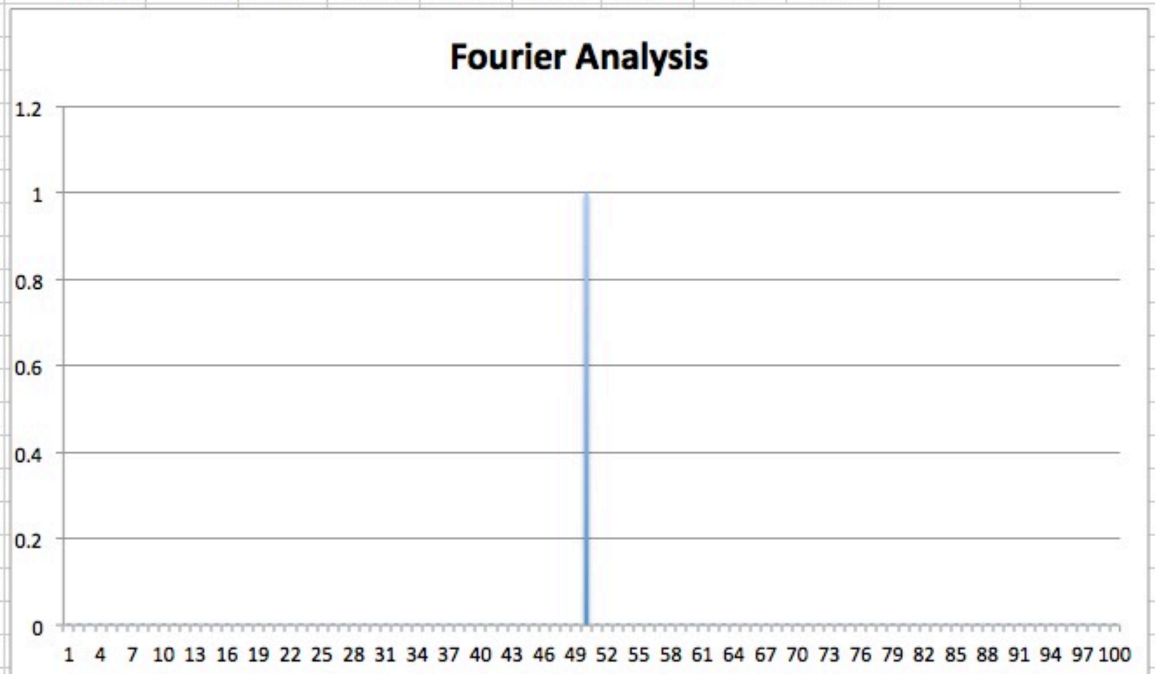

Recall that the phase does not affect the measurement:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	2.34	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0000											
48	0.0000											
49	0.0000											
50	1.0000											
51	0.0000											
52	0.0000											
53	0.0000											
		19	0.6956	0	0	0	0	0	0	0	0.6955589	

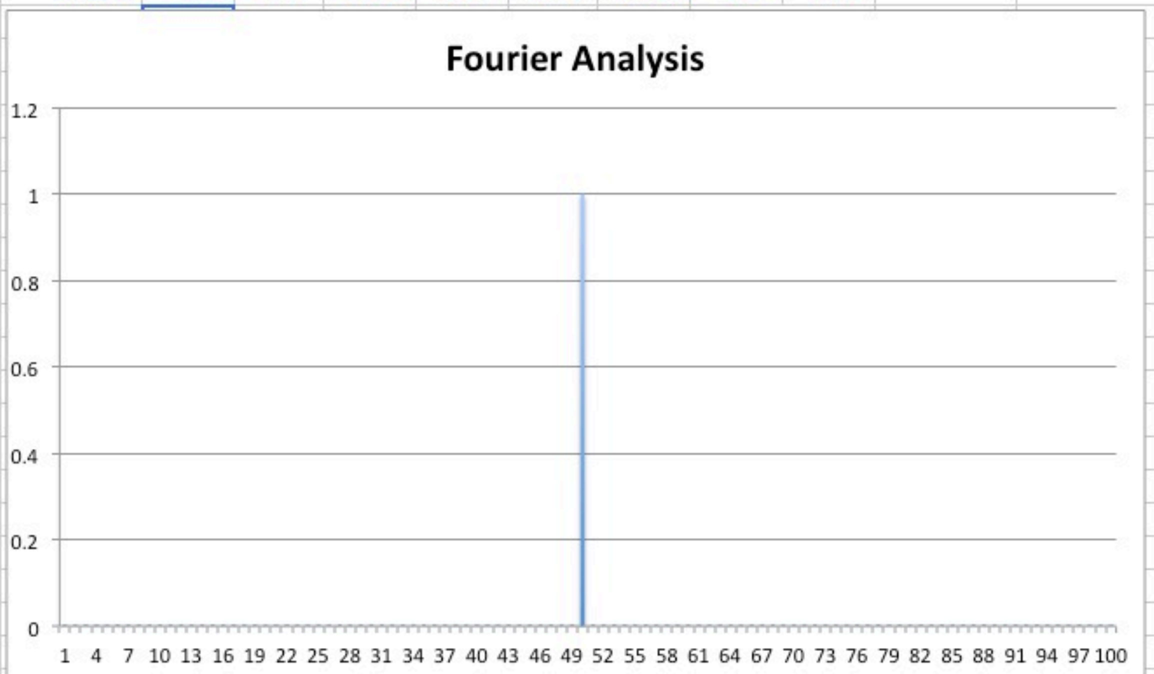

BOSTON UNIVERSITY
Computer Science

But let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								Signal	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave	
around 50 Hz		Freq:	50	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	Amplitude	Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	Measure											
47	0.0000											
48	0.0000											
49	0.0000											
50	1.0000											
51	0.0000											
52	0.0000											
53	0.0000											

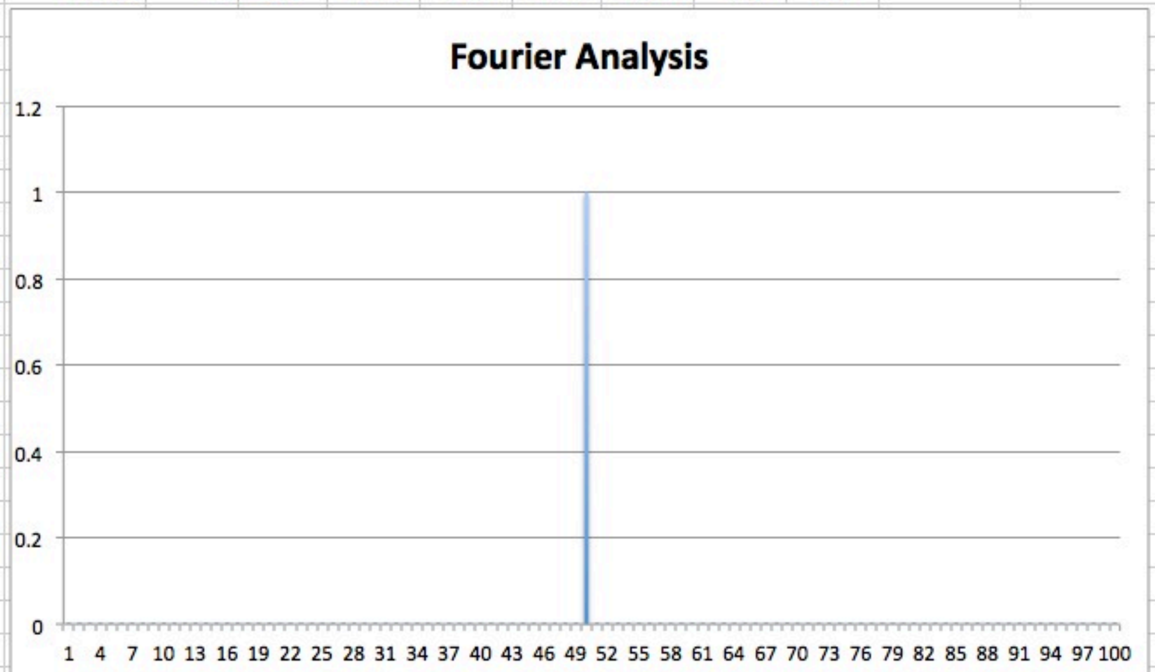

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

	Component Waves								Signal	
	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave	
Freq:	50.1	4	6	12	20	56.14	78	99		
Amp:	1	0	0	0	0	0	0	0	1	Freq:
Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		

Exact Measurements around 50 Hz

Probe Freq	Amplitude Measure
47	0.0309
48	0.0460
49	0.0885
50	0.9827
51	0.1102
52	0.0527
53	0.0349

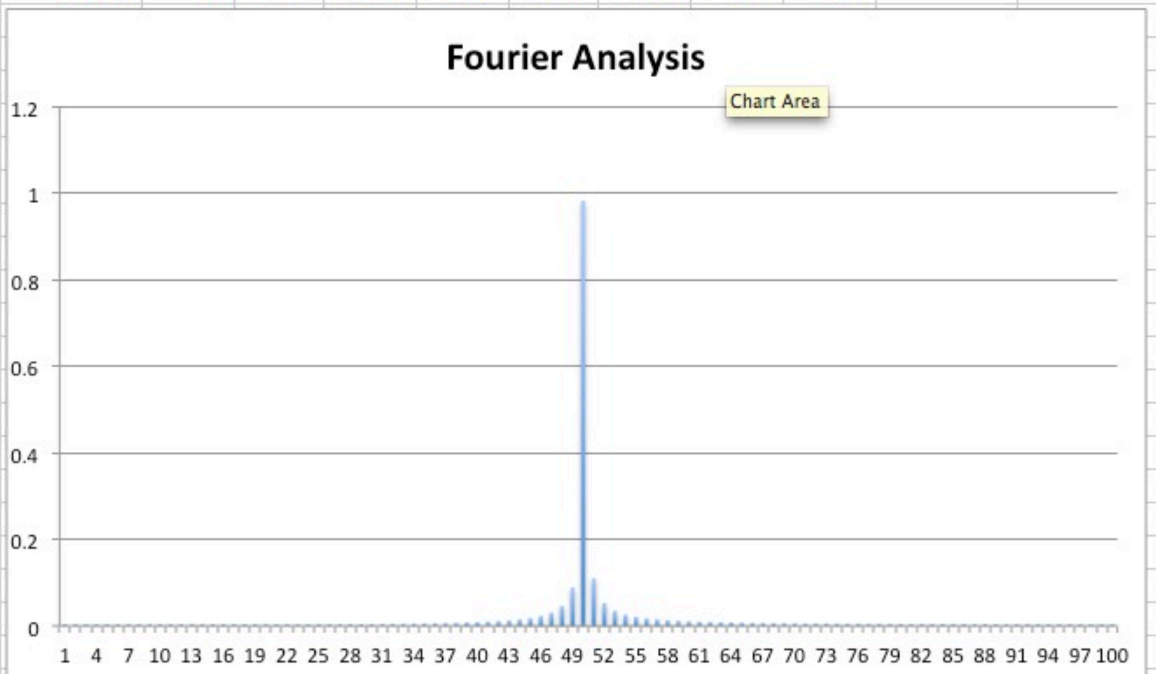

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50.2	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0557											
48	0.0823											
49	0.1531											
50	0.9327											
51	0.2367											
52	0.1068											
53	0.0696											
		19	-0.993	0	0	0	0	0	0	0	-0.992883	

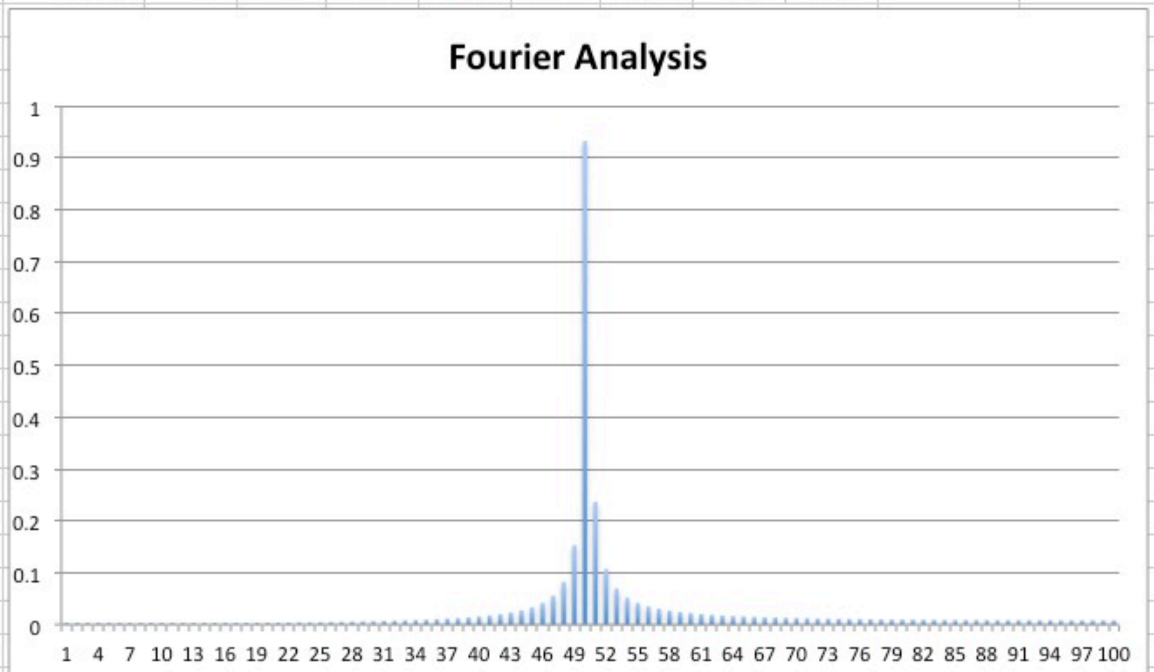

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								Signal	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave	
around 50 Hz		**Freq:**	50.3	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe Freq	**Amplitude Measure**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
47	0.0742											
48	0.1081											
49	0.1942											
50	0.8545											
51	0.3718											
52	0.1554											
53	0.0993											
		19	-0.984	0	0	0	0	0	0	0	-0.984011	

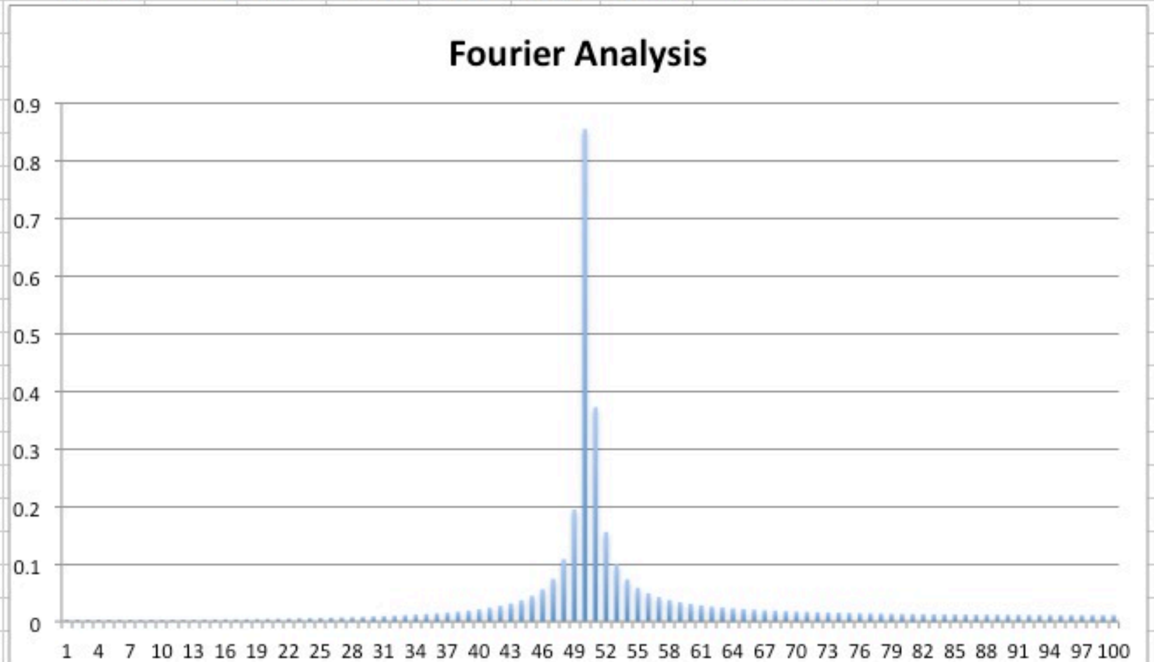

Computer Science

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50.4	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0863											
48	0.1234											
49	0.2134											
50	0.7540											
51	0.5074											
52	0.1921											
53	0.1194											
		19	-0.972	0	0	0	0	0	0	0	-0.971633	

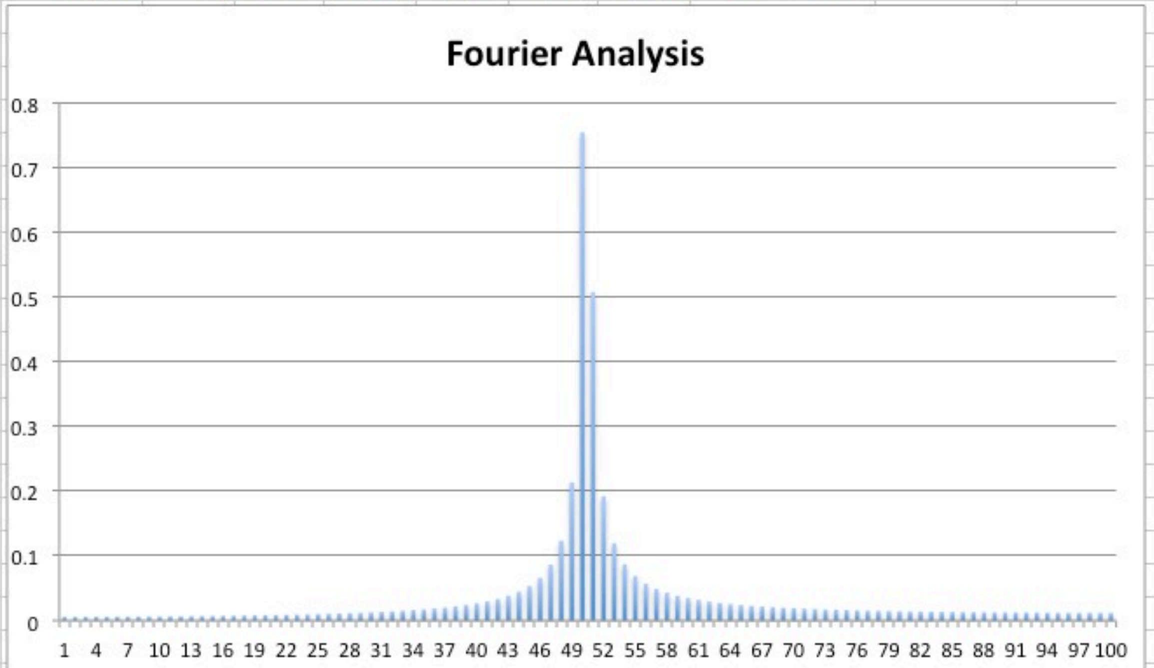

Digital Audio Fundamentals: The Discrete Fourier Transform

Boston University
Computer Science

Let’s see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Signal Sum Wave	
			Component Waves									
Exact Measurements												
around 50 Hz		**Freq:**	50.5	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
		Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		

Probe Freq	Amplitude Measure
47	0.0911
48	0.1274
49	0.2122
50	0.6366
51	0.6367
52	0.2124
53	0.1275

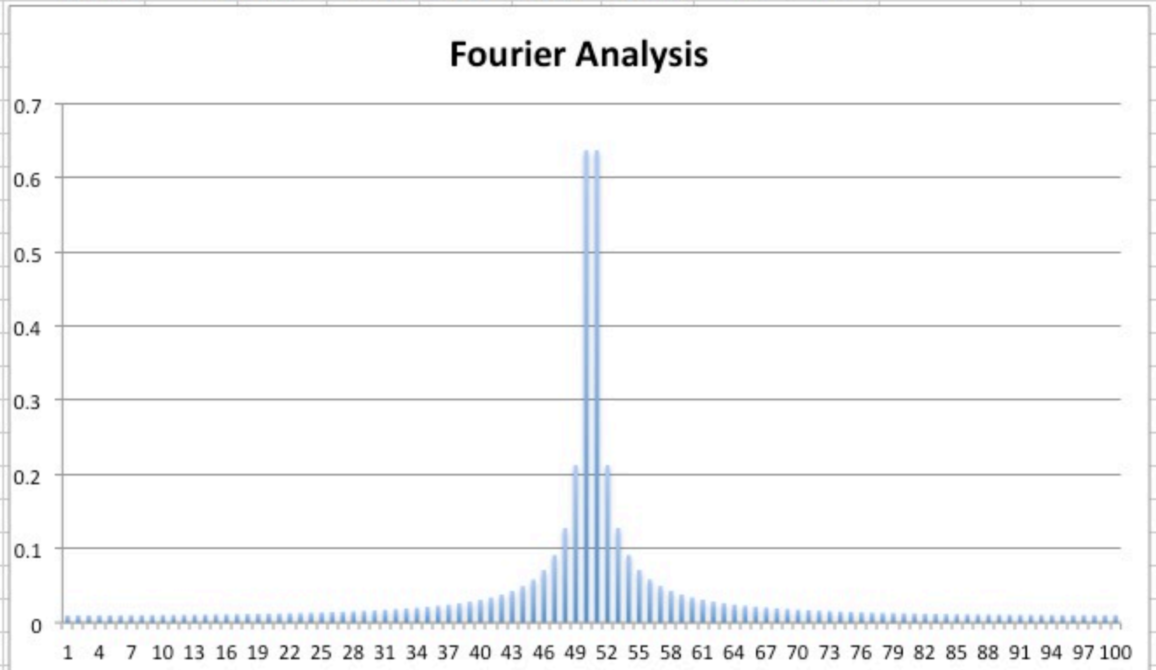

Digital Audio Fundamentals: The Discrete Fourier Transform

Computer Science

Let’s see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50.6	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0869											
48	0.1193											
49	0.1920											
50	0.5073											
51	0.7541											
52	0.2136											
53	0.1235											
		19	-0.937	0	0	0	0	0	0	0	-0.936552	

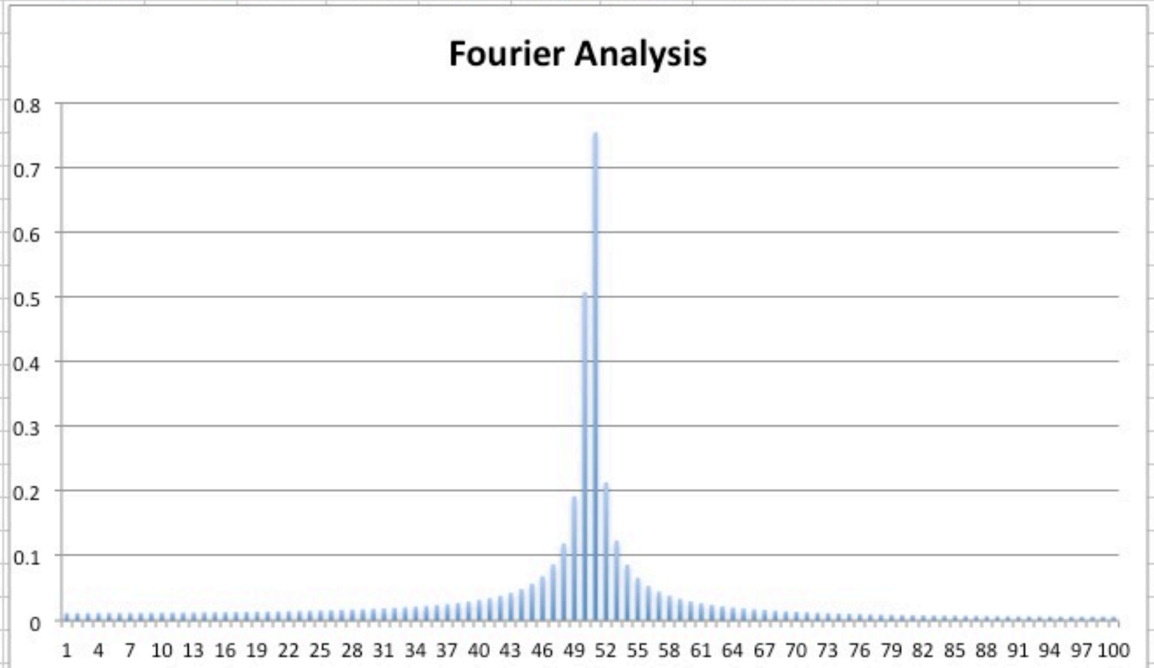

Let’s see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

		Component Waves								Signal	
		Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave	
Exact Measurements around 50 Hz	Freq:	50.7	4	6	12	20	56.14	78	99		
	Amp:	1	0	0	0	0	0	0	0	1	Freq:
	Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		

Probe Freq	Amplitude Measure
47	0.0735
48	0.0992
49	0.1553
50	0.3717
51	0.8546
52	0.1943
53	0.1082

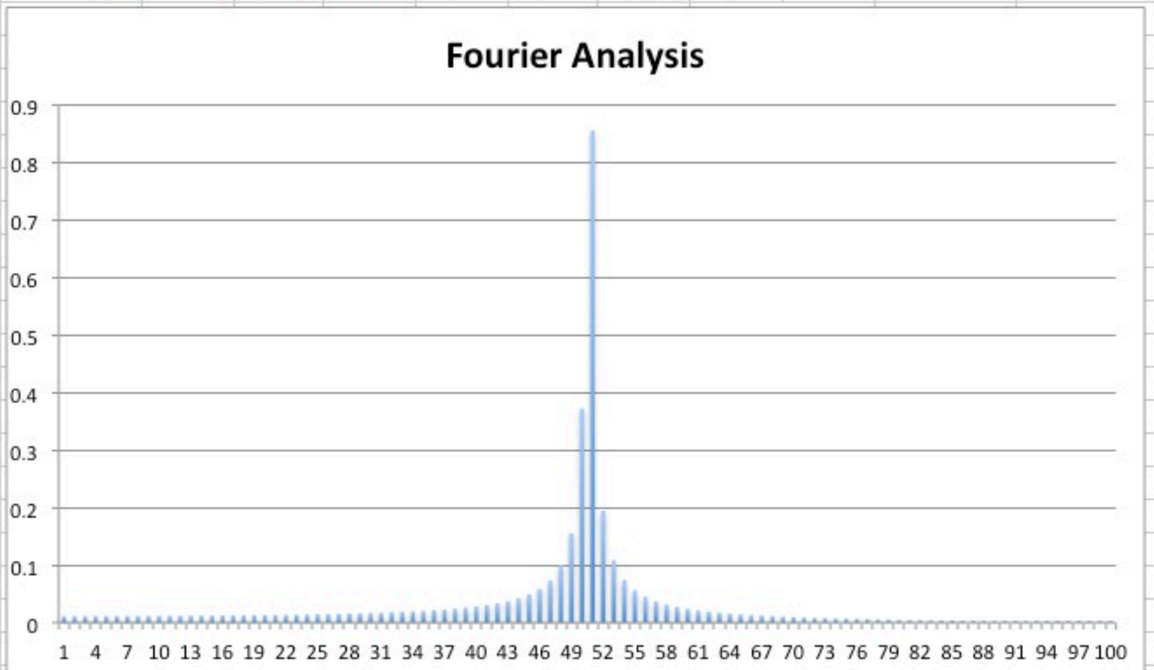

19 -0.914 0 0 0 0 0 0 0 -0.913973

Digital Audio Fundamentals: The Discrete Fourier Transform

Computer Science

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50.8	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0521											
48	0.0697											
49	0.1068											
50	0.2367											
51	0.9327											
52	0.1531											
53	0.0822											
		19	-0.888	0	0	0	0	0	0	0	-0.888139	

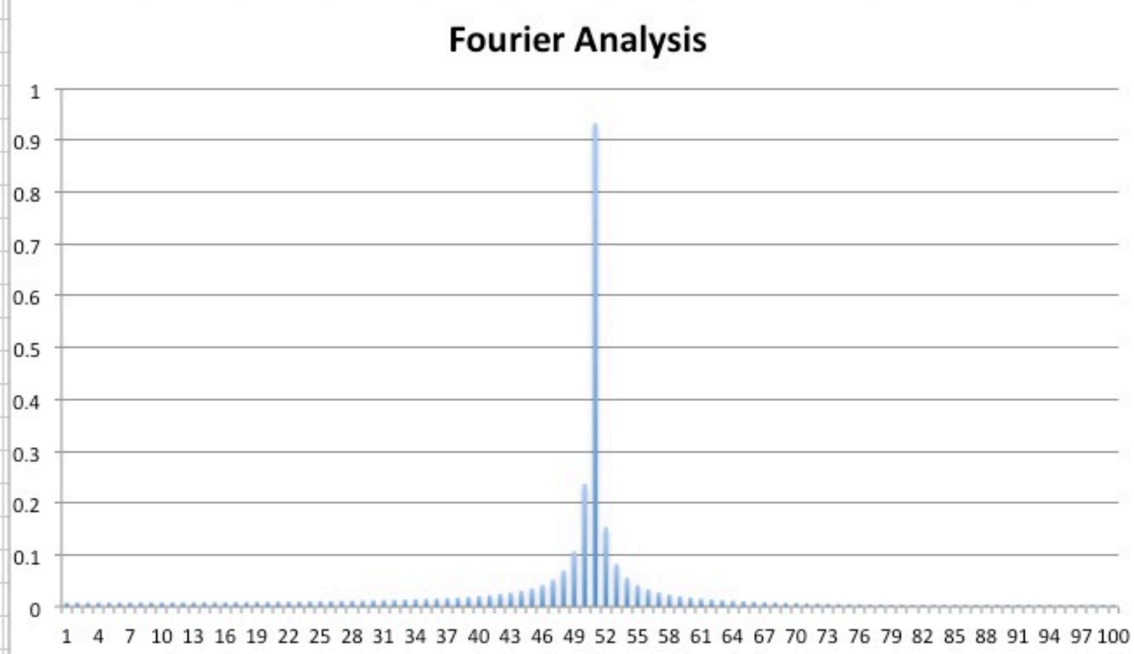

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

		Fourier Analysis on Input Signal Wave of 200 samples										
			Component Waves								**Signal**	
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	**Sum Wave**	
around 50 Hz		**Freq:**	50.9	4	6	12	20	56.14	78	99		
		Amp:	1	0	0	0	0	0	0	0	1	Freq:
Probe	**Amplitude**	**Phase:**	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
Freq	**Measure**											
47	0.0262											
48	0.0349											
49	0.0527											
50	0.1102											
51	0.9827											
52	0.0885											
53	0.0459											
		19	-0.859	0	0	0	0	0	0	0	-0.859142	

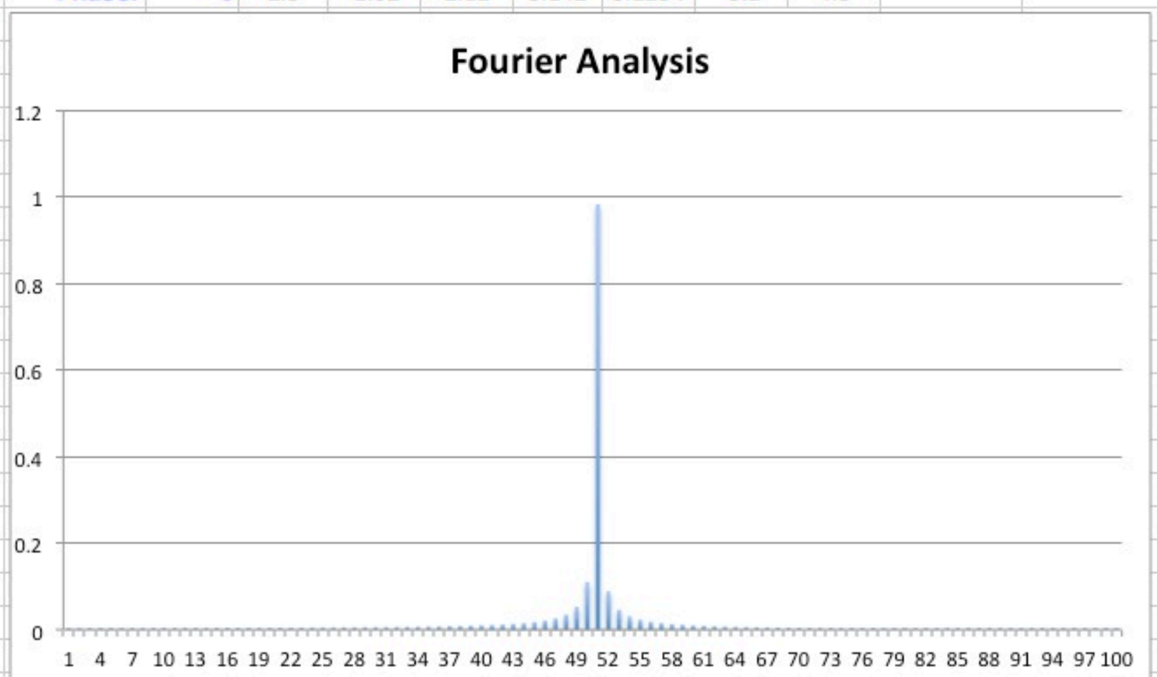

Digital Audio Fundamentals: The Discrete Fourier Transform

Computer Science

Let's see what happens as we change the frequency slowly from 50 to 51 Hz:

Fourier Analysis on Input Signal Wave of 200 samples

		Component Waves								Signal	
Exact Measurements		Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Sum Wave	
around 50 Hz	Freq:	51	4	6	12	20	56.14	78	99		
	Amp:	1	0	0	0	0	0	0	0	1	Freq:
	Phase:	0	2.3	-1.02	2.12	3.142	3.1234	-3.2	4.3		
	19	-0.827	0	0	0	0	0	0	0	-0.827084	

Probe Freq	Amplitude Measure
47	0.0000
48	0.0000
49	0.0000
50	0.0000
51	1.0000
52	0.0000
53	0.0000

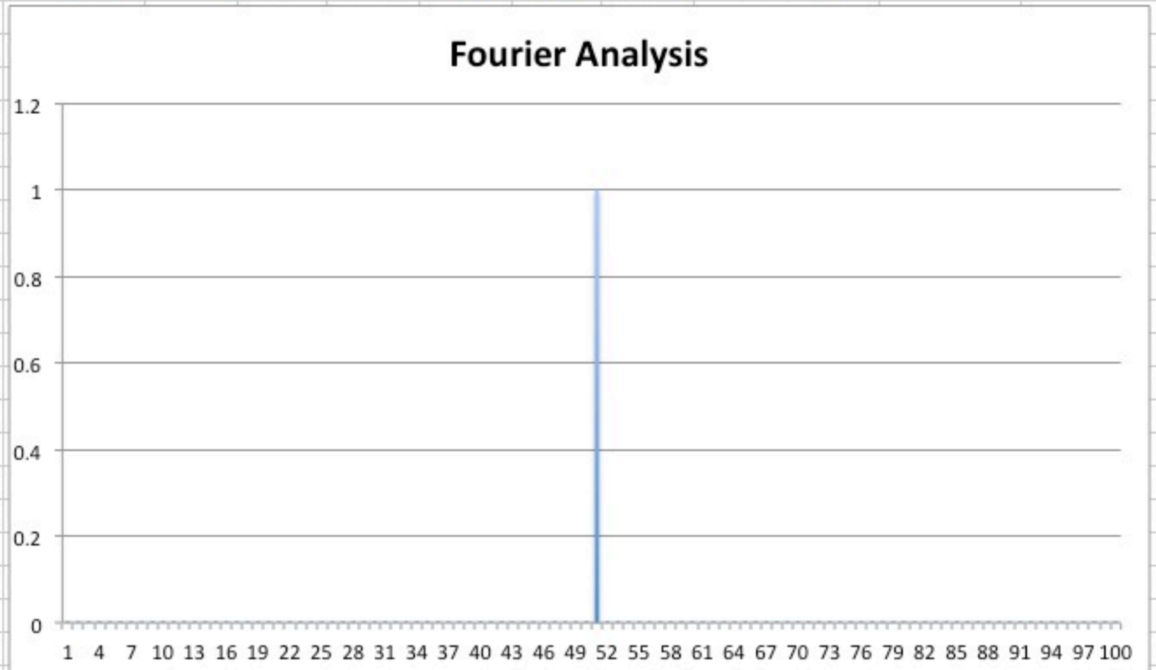

You can see this in a typical spectrum, where the characteristic shape of a frequency component ("triangular mountain") shows up repeatedly:

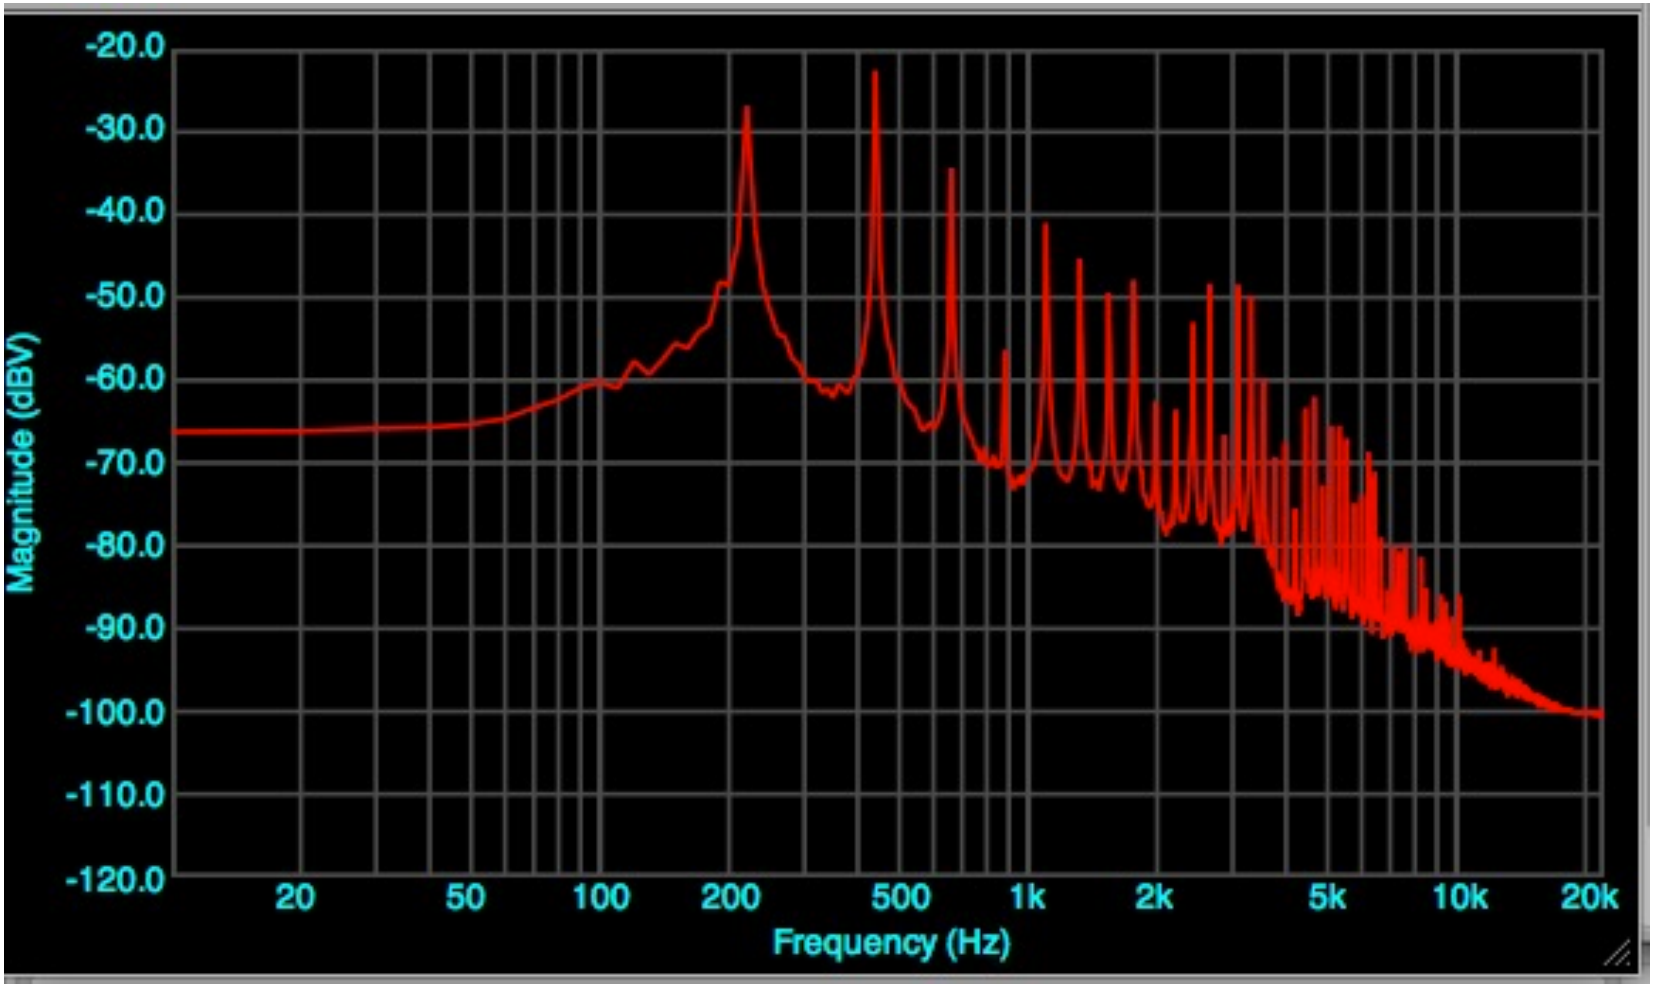

The typical **solution** used is to de-emphasize the signal components at the edges, by tapering the amplitude of the signal using either a triangular function (which is used to modify the amplitude of the signal)

W(n,N) = 1 – Abs[(n – (N/2))/(N/2)]   for 1 <= n <= N

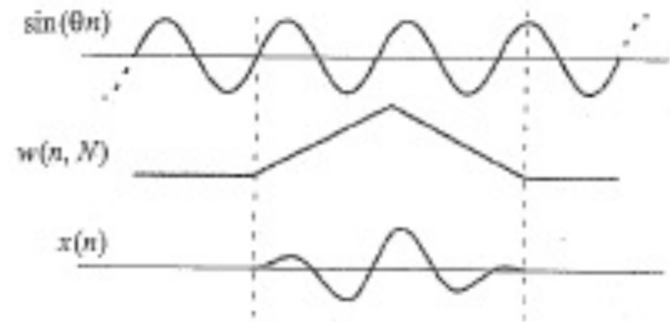

Figure 3.19
Windowing with a triangular function.

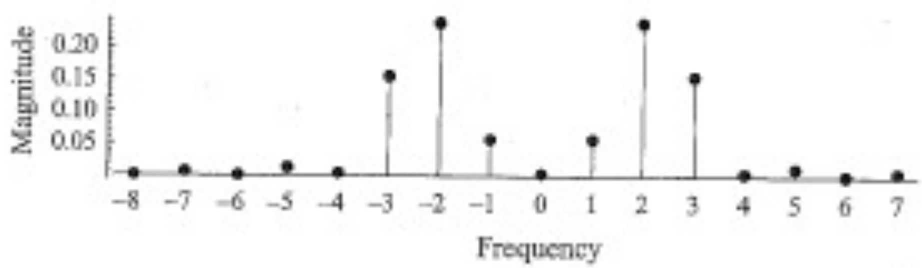

Figure 3.20

Note: We can expect that this approach will change the amplitude measurement, since it reduces the overall sum of the samples!

Or some more complex function:

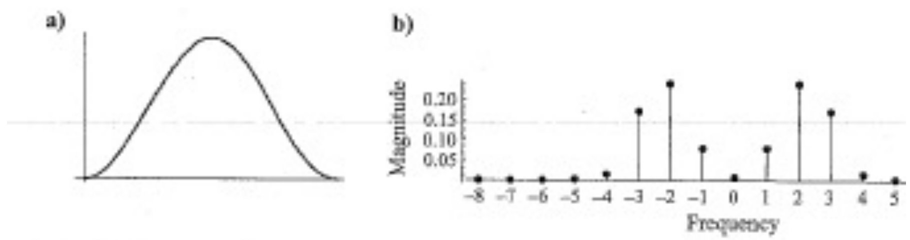

Figure 3.21
Hann window and resulting spectrum.

$$H(x, n, N) = \begin{cases} (1-a)\cos\left(2\pi\frac{n}{N}+\pi\right)+a, & 0 \le n < N, \\ 0 & \text{otherwise,} \end{cases}$$

where $a = 1/2$. It is shown in figure 3.21a.

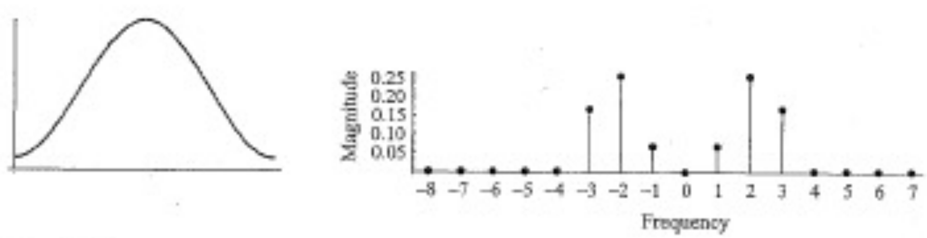

Figure 3.22
Hamming window and resulting spectrum.

A nice description of the various ”window functions,” with interesting graphics, is provided by http://en.wikipedia.org/wiki/Window_function

Computer Science

Let's try two of these in our experiment on a non-integral frequency of 50.4 Hz, using the Triangular and the Hann Windows:

Windowing Functions

Window size N = 200

Samp Num	Triangular	Hann
0	0	2.14617E-09
1	0.01	0.000245252
2	0.02	0.000983671
3	0.03	0.002214529
4	0.04	0.003936613
5	0.05	0.006148223
6	0.06	0.008847175
7	0.07	0.012030808
8	0.08	0.015695979
9	0.09	0.019839072
10	0.1	0.024455997
11	0.11	0.029542199
12	0.12	0.035092659
13	0.13	0.041101899
14	0.14	0.04756399
15	0.15	0.054472554
16	0.16	0.061820774
17	0.17	0.069601398
18	0.18	0.077806748
19	0.19	0.086428728
20	0.2	0.095458829
21	0.21	0.10488814
22	0.22	0.114707355
23	0.23	0.124906785
24	0.24	0.135476365
25	0.25	0.146405664

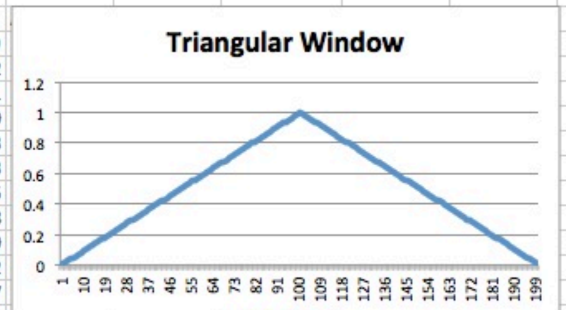

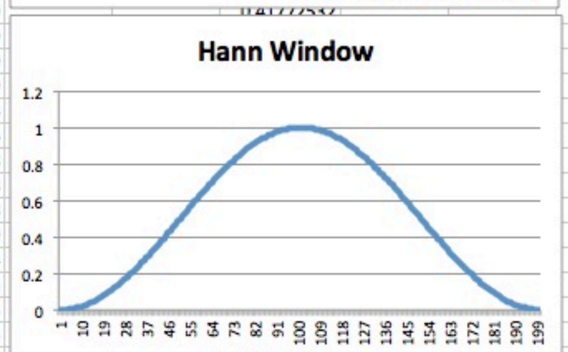

Let's try two of these in our experiment on a non-integral frequency of 50.4 Hz, using the Rectangular (as before), the Triangular, and the Hann Windows:

Fourier Analysis on Input Signal Wave of 200 samples

	Component Waves											Signal
	Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Windowing Functions			Sum Wave
Freq:	50.4	4	6	12	20	56.14	78	99	Rect	Triang	Hann	
Amp:	1	0	0	0	0	0	0	0	1	0	0	2
Phase:	3	2.3	-1.02	2.12	3.1415	3.1234	-3.2	4.3				
Samp Num							Num Samps:					200

Probe Waves of frequencies 1, 2, 3, ..., 100

Freq:	1	2	3	4

Exact Measurements around 50 Hz

Probe Freq	Amplitude Measure
47	0.0853
48	0.1224
49	0.2125
50	0.7530
51	0.5084
52	0.1930
53	0.1203

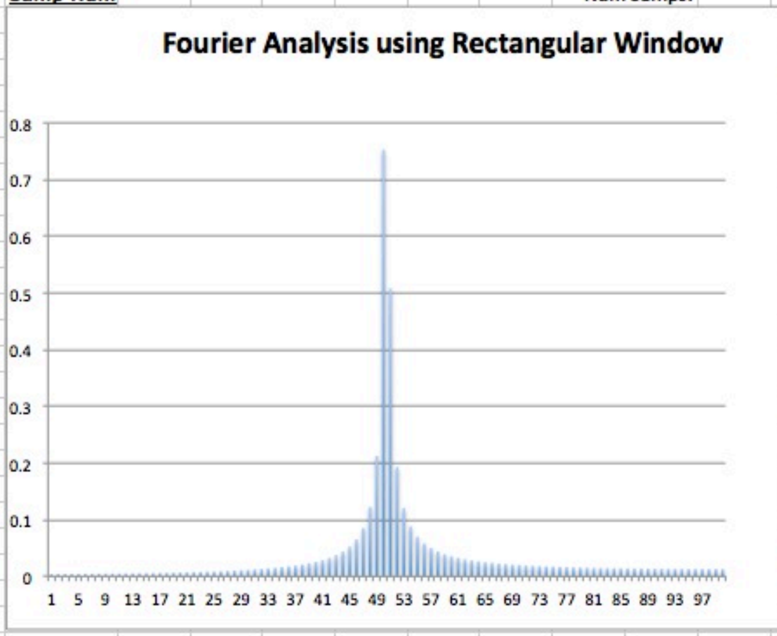

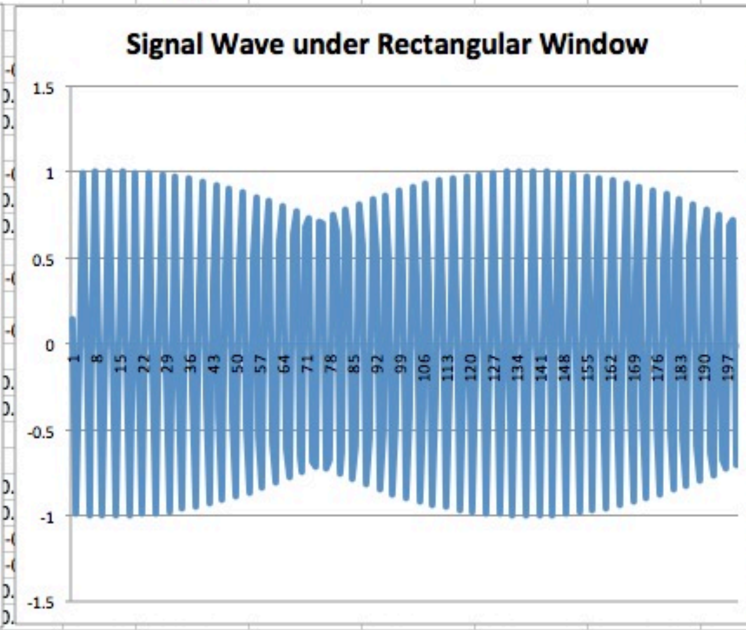

Let’s try two of these in our experiment on a non-integral frequency of 50.4 Hz, using the Rectangular (as before), the Triangular, and the Hann Windows:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves										Signal			
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Windowing Functions		Sum Wave		Probe Waves of frequencies 1, 2, 3, ..., 100	
around 50 Hz		Freq:	50.4	4	6	12	20	56.14	78	99	Rect	Triang	Hann			
		Amp:	1	0	0	0	0	0	0	0	0	1	0	2	Freq:	1 2 3 4
Probe	Amplitude	Phase:	3	2.3	-1.02	2.12	3.1415	3.1234	-3.2	4.3						
Freq	Measure	Samp Num							Num Samps:				200			

Probe Freq	Amplitude Measure
47	0.0115
48	0.0122
49	0.0677
50	0.4376
51	0.3685
52	0.0274
53	0.0197

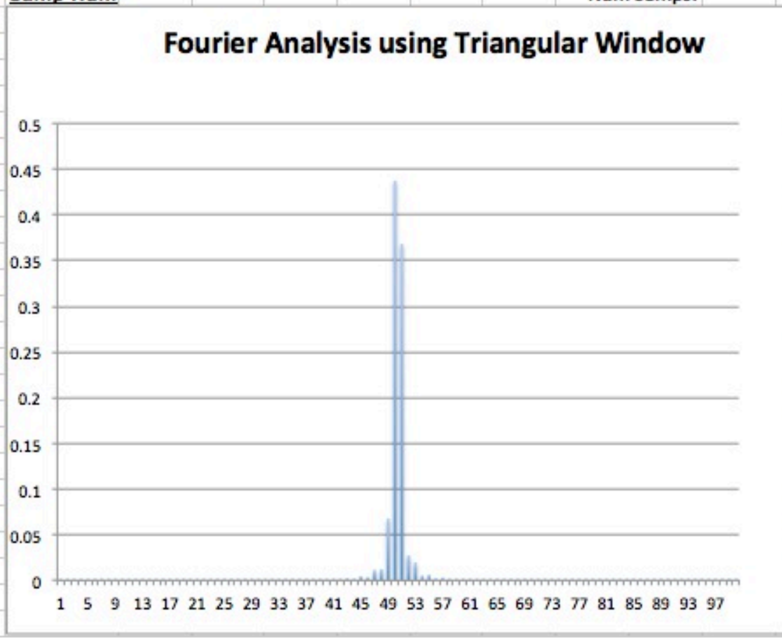

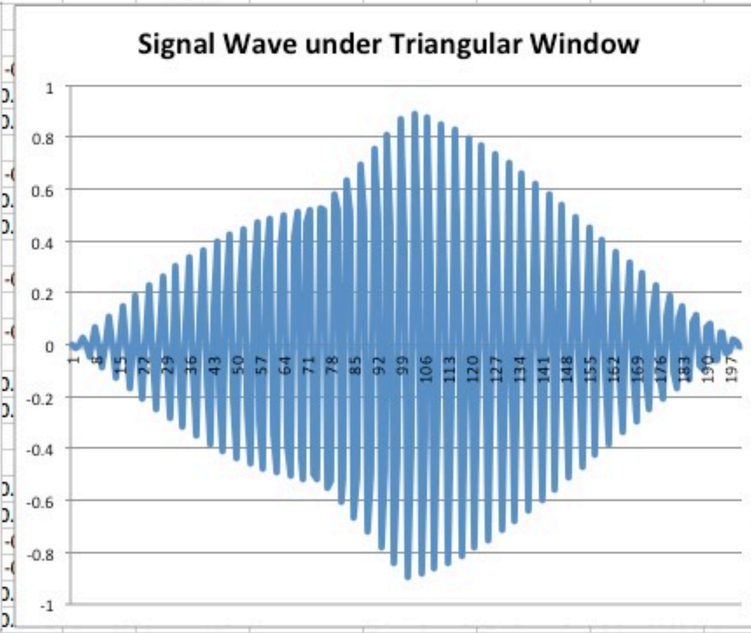

0
33321
58984
12448
75358
78515
54699
51313
32781
32695
05644
22872
02671
02675
28732
05664
32723
32816
51354
54747
78568
75416
12509
59047

Let’s try two of these in our experiment on a non-integral frequency of 50.4 Hz, using the Rectangular (as before), the Triangular, and the Hann Windows:

Fourier Analysis on Input Signal Wave of 200 samples

			Component Waves											Signal				
Exact Measurements			Wave 1	Wave 2	Wave 3	Wave 4	Wave 5	Wave 6	Wave 7	Wave 8	Windowing Functions			Sum Wave		Probe Waves of frequencies 1, 2, 3, ..., 100		
around 50 Hz		Freq:	50.4	4	6	12	20	56.14	78	99	Rect	Triang	Hann					
		Amp:	1	0	0	0	0	0	0	0	0	0	1	2	Freq:	1	2	3	4
Probe	Amplitude	Phase:	3	2.3	-1.02	2.12	3.1415	3.1234	-3.2	4.3								
Freq	Measure	Samp Num							Num Samps:					200				
47	0.0042																	
48	0.0132																	
49	0.1126																	
50	0.4505																	
51	0.3942																	
52	0.0606																	
53	0.0101																	

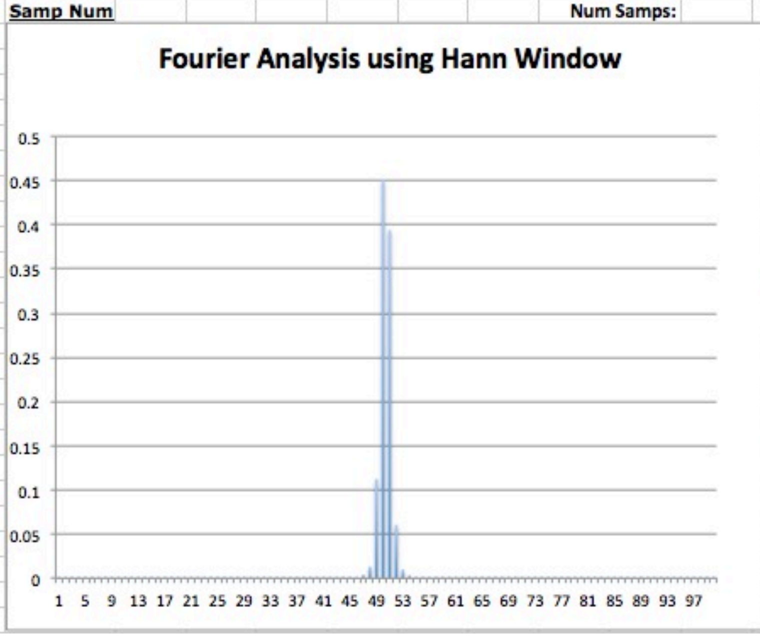

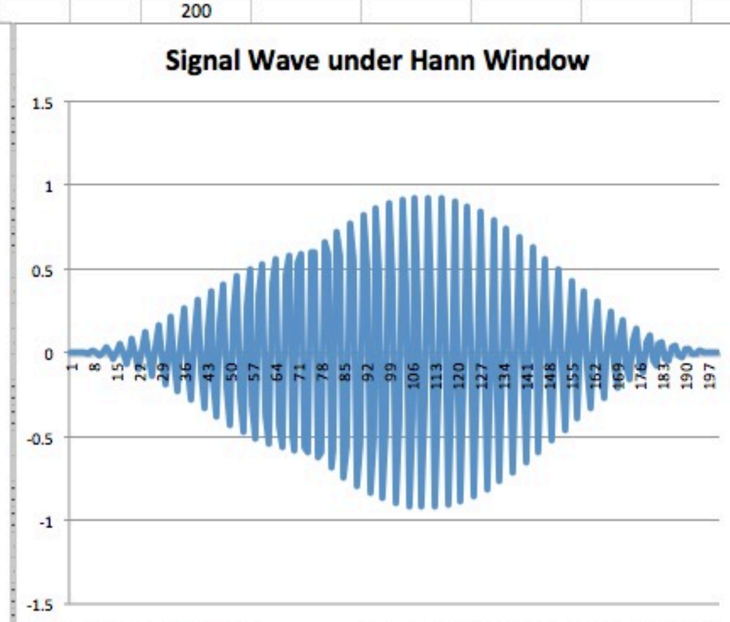

0
33321
68984
12448
.75358
78515
54699
51313
32781
82695
.05644
22872
02671
02675
28732
.05664
82723
32816
51354
54747
78568
.75416
12509
69047

Conclusions on windowing for the DFT:

(1) Window size determines frequency resolution: given a window size of W samples, with a fundamental frequency of f = SR / W, we can only probe for the integral frequencies (the harmonics of F):

0, f, 2*f, 3*f, …., k*f, ….., Nyquist Limit

Any other frequencies will be subject to the "picket fence" problem and only approximated.

(2) Non-integral frequencies cause "leakage" to adjacent integral frequencies; good windowing functions (e.g., Hann) mitigate leakage effects and provide reasonably accurate measurements of amplitude of components, after correction.

Professional tools such as Electroacoustics Toolbox allow you to set these features, as well as window length, whether windows overlap, whether and how to average the successive measurements, whether and how to weight the measures to the psychoacoustical properties of human hearing, how to display the result, etc., etc., etc. and to output the analysis to a file.

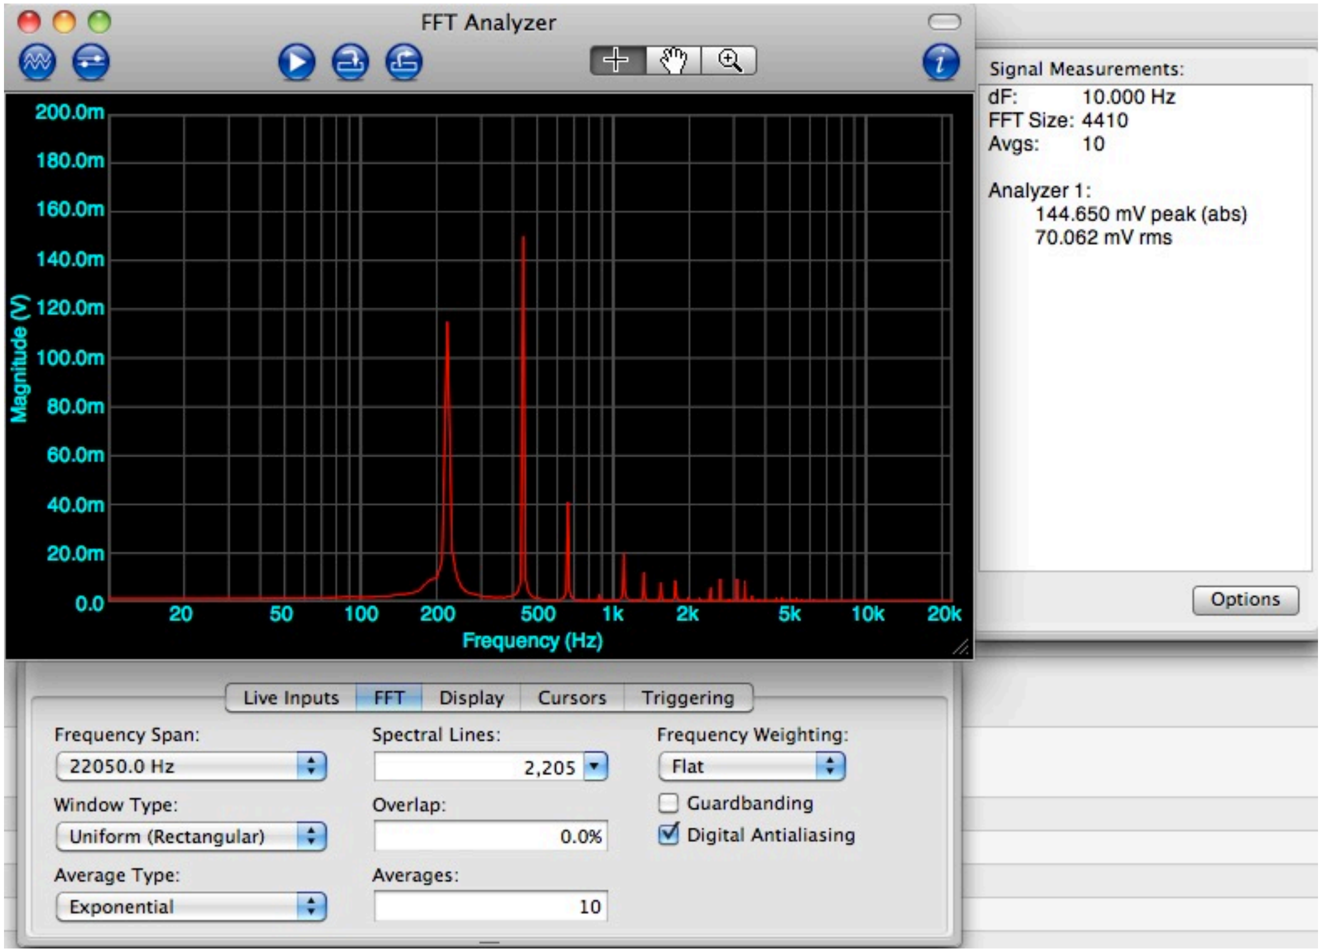

Digital Audio Fundamentals: The Discrete Fourier Transform

Computer Science

Professional tools such as Electroacoustics Toolbox allow you to set these features, as well as window length, whether windows overlap, whether and how to average the successive measurements, whether and how to weight the measures to the psychoacoustical properties of human hearing, how to display the result, etc., etc., etc. and to output the analysis to a file.

```
Created by Electroacoustics Toolbox, version 3.0.1 (rev. 1109), on Thursd
Project Name:   untitled
Module Type:    FFT Analyzer
Module Name:    FFT Analyzer

Device: Soundflower (2ch)
Device Unique ID:       SoundflowerEngine:0
Channel Label:  Analyzer 1
Units:  Vpk
Sensitivity:    1.00000000 V/V
Frequency Weighting:    Flat
Sample Rate:    44.10000000 kHz
Sample Period:  22.67573655 µs
Decimation Factor:      1
Digital Antialiasing:   ON
FFT Length:     4410
Guardbanding:   OFF
Spectral Lines: 2206
Frequency Resolution:   10.00000000 Hz
FFT Time Record Length: 100.00000149 ms
Window: Uniform (Rect)
Averaging:      Exponential
Averages:       10
Total Time Record Length:       1.00000000 s
Overlap:        0.00000000 %

Frequency       Magnitude
0       0.0009345763
10      0.001883723
20      0.00192472
```

220	0.2791846
230	0.04434038
240	0.01913349
250	0.01382732
260	0.01114911
270	0.00882916
280	0.007373199
290	0.0065718
300	0.006028267
310	0.005524478
320	0.0049454
330	0.004598976
340	0.004576757
350	0.004409567
360	0.004563091
370	0.004502238
380	0.004490616
390	0.004606624
400	0.004957298
410	0.005889128
420	0.007712504
430	0.01238511
440	0.1653076
450	0.01271653
460	0.005566336
470	0.003511857
480	0.002436093
490	0.002428316