EE3S1 Signal Processing – DSP

Lecture 5: Spectral analysis (Ch. 14)

Alle-Jan van der Veen

–

26 December 2025

Contents

How do we measure the spectral content of a *finite* segment of a signal?

What is the effect of sampling in frequency?

- DFT of a sinusoid [Ch. 10.2.5]
- Resolution [Ch. 10.2.6]
- Zero padding [Ch. 10.5.1]

- Windowing [Ch. 14.1]
- Short-time Fourier transform (STFT) [Ch. 14.2]
- Periodogram [Ch. 14.3, only as application]

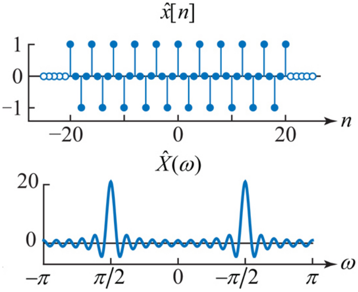

Recap - DFT

The DFT $X[k]$ is obtained by sampling the DTFT $X(\omega)$ of a length-N sequence:

$$\text{DTFT:} \quad X(\omega) = \sum_{n=0}^{N-1} x[n] e^{-j\omega n}$$

$$\text{DFT:} \quad X[k] = X\left(\frac{2\pi}{N} k\right) = \sum_{n=0}^{N-1} x[n] e^{-j\frac{2\pi}{N}kn}$$

$$k = 0, \cdots, N-1$$

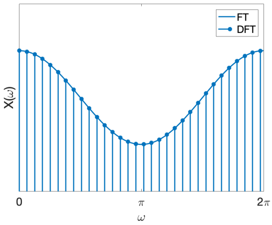

N samples $x[n]$ in time are mapped to N samples $X[k]$ in frequency.

$$\text{IDFT:} \quad x[n] = \frac{1}{N} \sum_{n=0}^{N-1} X[k] e^{j\frac{2\pi}{N}kn}, \quad n = 0, \cdots, N-1$$

(Sampling gives rise to aliasing and periodicity, not the topic for today.)

Recap - DFT

Reconstruction: $X(\omega)$ can be found back using interpolation:

$$X(\omega) = \frac{1}{N} \sum_{k=0}^{N-1} X[k] G\left(\omega - \frac{2\pi}{N} k\right)$$

where $G(\omega)$ is the Dirichlet kernel, a "periodic sinc" function:

$$G(\omega) = \frac{\sin(\frac{1}{2}\omega N)}{\sin(\frac{1}{2}\omega)} e^{-j\omega \frac{N-1}{2}}$$

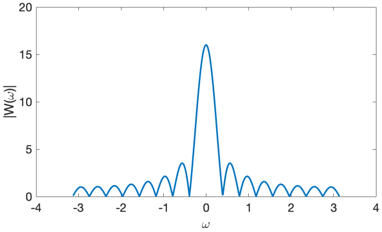

Resolution of the DFT [Ch. 10.2.6]

Suppose we sample a signal $x(t)$ with rate $F_s = \frac{1}{T_s}$, and collect N samples $x[k]$. The total time span is $T = N\,T_s = \frac{N}{F_s}$.

The DFT gives us N samples $X[k]$, spaced uniformly between 0 and 2π on the ω-axis.

$$\Delta\omega = \frac{2\pi}{N} \text{ [rad]}$$

$$\Delta F = \frac{F_s}{N} \text{ [Hz]}$$

Hence $\Delta F = \frac{1}{T}$!

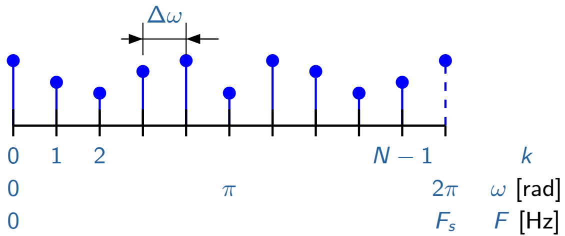

- Only depends on the total duration of the signal, the "aperture"
- Does not depend on the sample rate or number of samples!

Zero padding [Ch. 10.5.1]

Suppose $x[n]$ is zero except for $n = 0, \cdots, L-1$, and its DTFT is $X(\omega)$.

- A DFT on the L samples gives $X[k]$ (L samples), and we could use interpolation with the Dirichlet kernel to find $X(\omega)$ for any ω.
- But if our aim is to have more samples in frequency domain for a nicer plot (larger $N \Rightarrow$ smaller spacing $2\pi/N \Rightarrow$ smoother plot), it is more convenient to use zero padding:

For any $N \geq L$, the DFT of $x[n]$ (i.e., the L nonzero samples, extended with $N - L$ zero samples) gives $X[k]$, $k = 0, \cdots, N-1$, which are samples of $X(\omega)$.

Proof: Essentially, this follows from

$$X(\omega) = \sum_{n=0}^{L-1} x[n]e^{-j\omega n} = \sum_{n=0}^{N-1} x[n]e^{-j\omega n}$$

Zero padding

Example ($L = 3$)

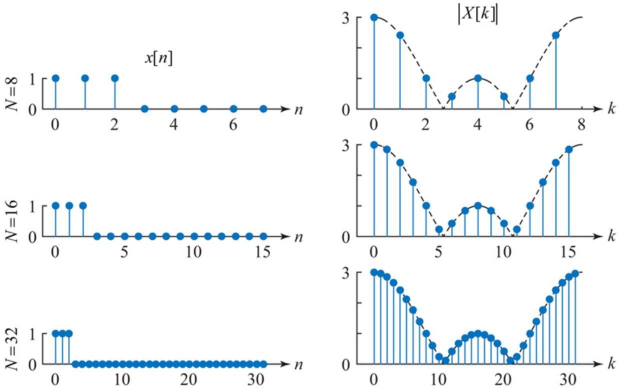

Zero padding doesn't give extra information, but improves visualization.

The actual resolution is not determined by N but by L (see later).

DFT of a sinusoid [Ch. 10.2.5]

- Complex exponential at a DFT frequency: $x[n] = e^{j\frac{2\pi}{N}\ell n}$

$$X[k] = \sum_{n=0}^{N-1} e^{j\frac{2\pi}{N}\ell n} e^{-j\frac{2\pi}{N}kn} = \sum_{n=0}^{N-1} e^{j\frac{2\pi}{N}(\ell-k)n} = N\delta[(k-\ell)_N]$$

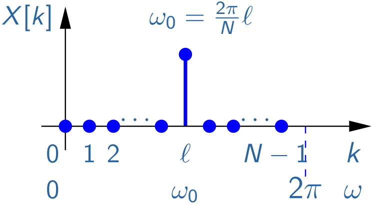

- Harmonic function: $x[n] = e^{j(\omega_0 n + \phi_0)}$ results in

$$X[k] = e^{j\phi_0} \sum_{n=0}^{N-1} e^{j(\omega_0 - \frac{2\pi}{N}k)n} = \begin{cases} N\, e^{j\phi_0}\delta(\omega_0 - \frac{2\pi}{N}k) & \text{if } \omega_0 = \frac{2\pi}{N}\ell \\ e^{j\phi_0}\dfrac{1 - e^{j(N\omega_0 - 2\pi k)}}{1 - e^{j(\omega_0 - \frac{2\pi}{N}k)}} & \text{else} \end{cases}$$

DFT of a sinusoid

Example 1

$$x[n] = \sin(\omega_0 n), \quad 0 \le n \le N-1$$

where ω_0 is an integer multiple of $2\pi/N$
e.g. $\omega_0 = 2\pi \cdot 5/N$ with $N = 60$:

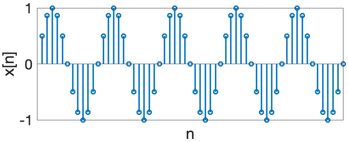

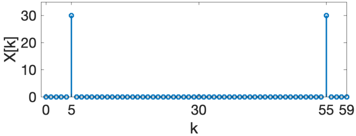

- Exactly an integer number of periods of the sinusoid are sampled.
- $X(\omega)$ obtained via zero padding (red line) shows Dirichet functions centered at the spikes

DFT of a sinusoid

Example 1

$$x[n] = \sin(\omega_0 n), \quad 0 \leq n \leq N-1$$

where ω_0 is an integer multiple of $2\pi/N$
e.g. $\omega_0 = 2\pi \cdot 5/N$ with $N = 60$:

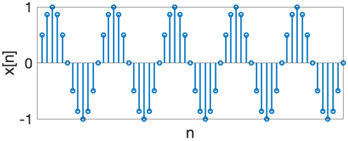

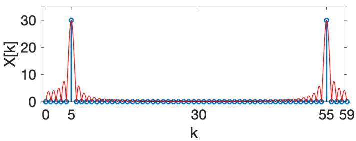

- Exactly an integer number of periods of the sinusoid are sampled.
- $X(\omega)$ obtained via zero padding (red line) shows Dirichet functions centered at the spikes

DFT of a sinusoid

Example 2

$$x[n] = \sin(\omega_0 n), \quad 0 \leq n \leq N-1$$

where ω_0 is not an integer multiple of $2\pi/N$
e.g. $\omega_0 = 2\pi \cdot 5.5/N$ with $N = 60$:

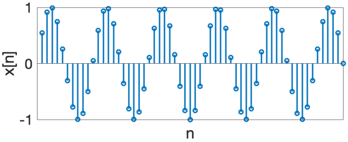

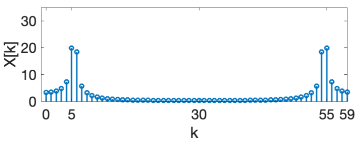

- In this case, ω_0 falls in between two sample points. The peak could be localized more accurately using zero padding.
- The DFT can be used to estimate a frequency, but its resolution is not great for small N.

DFT of a sinusoid

Example 2

$$x[n] = \sin(\omega_0 n), \quad 0 \le n \le N-1$$

where ω_0 is not an integer multiple of $2\pi/N$
e.g. $\omega_0 = 2\pi \cdot 5.5/N$ with $N = 60$:

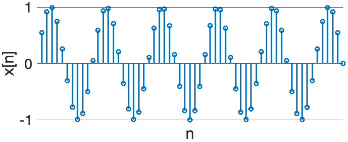

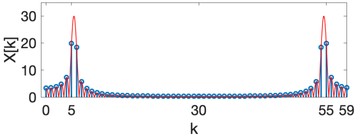

- In this case, ω_0 falls in between two sample points. The peak could be localized more accurately using zero padding.
- The DFT can be used to estimate a frequency, but its resolution is not great for small N.

Spectral analysis using the DFT [Ch. 14.1]

We now look in more detail at the construction of spectra.

- For a signal $x[n]$, we would like to find

$$X(\omega) \equiv \sum_{n=-\infty}^{\infty} x[n] e^{-j\omega n}$$

  We need all the samples over an infinite interval.

- However, in practice, we only have N samples of $x[n]$ i.e., $\hat{x}[n]$ with $n = 0, 1, \cdots, N-1$.
  An estimate of $X(\omega)$ is

$$\hat{X}(\omega) = \sum_{n=0}^{N-1} x[n] e^{-j\omega n}$$

  (In practice: DFT with zero padding to approximate a continuous ω)

How good is this estimate? And can we improve on it?

Spectral analysis: general considerations

$\hat{x}[n]$ can be viewed as a windowed version of $x[n]$:

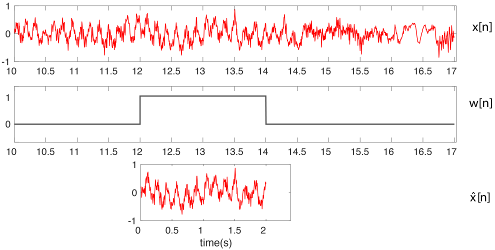

$$\hat{x}[n] = x[n]w[n]$$

Spectral analysis: general considerations

Due to the windowing, the spectrum of the finite sequence is expressed as a convolution of the spectrum of the original sequence and the Fourier transform of the window sequence:

$$\hat{x}[n] = x[n]w[n] \quad \Leftrightarrow \quad \hat{X}(\omega) = \frac{1}{2\pi}\int_{-\pi}^{\pi} X(\theta)W(\omega-\theta)\mathrm{d}\theta$$

For the rectangular window:

$$W(\omega) = \sum_{n=0}^{N-1} e^{-j\omega n} = \frac{1-e^{-j\omega N}}{1-e^{-j\omega}}$$

$$= \frac{\sin(\omega N/2)}{\sin(\omega/2)} e^{-j\omega(N-1)/2}$$

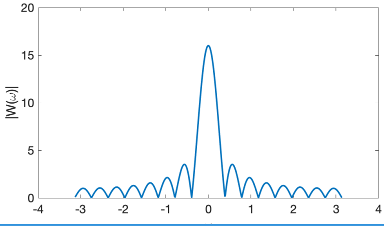

Interpretation

Convolution in frequency domain by the Dirichlet kernel broadens spectral lines in $X(\omega)$.

Assume $x[n]$ is a sinusoid:

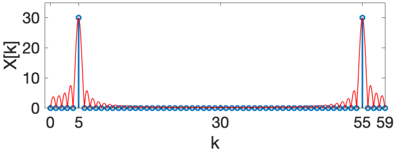

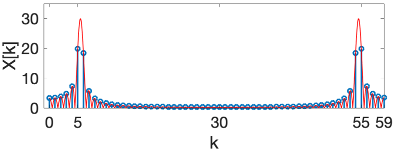

The red line is $\hat{X}(\omega)$ and is obtained by zero padding.

- The spectal line is broadened, limiting the resolution, i.e. our ability to distinguish closely spaced frequencies.
- The windowed spectrum is spread out over the whole frequency range – "spectral leakage".

Rectangular window

$$W(\omega) = \frac{\sin(\omega N/2)}{\sin(\omega/2)} e^{-j\omega(N-1)/2}$$

- Main lobe has a width of $\Delta\omega = \frac{4\pi}{N}$ (distance between two zero crossings)
- Sidelobes have an amplitude of -13 dB.

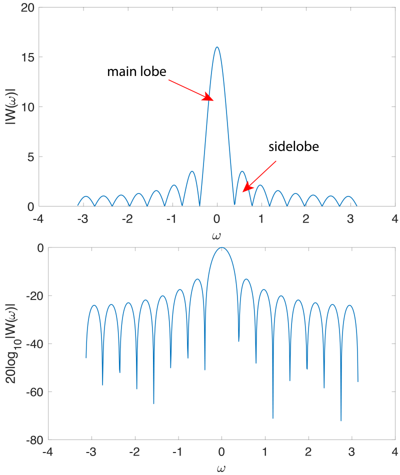

$N = 16$

Effect of windowing

- **Spectral smoothing:** Due to the non-zero width of the main lobe, two closely spaced peaks in the Fourier spectum may appear as a single peak in the DFT of the finite sequence.
  This relates to *resolution*. To distinguish two closely spaced frequencies, they need to be separated by more than the width of the main lobe.

- **Spectral leakage:** The spectrum is spread out to the whole frequency range. A weak peak in the original spectrum may be masked by the “leakage” from a large peak.
  This relates to *contrast*.

Effect of windowing - spectral smoothing

$x[n] = 1/3(\cos(\omega_1 n) + \cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, with $n = 0, 1, \cdots N-1$, $N = 100$ and $\omega_2 = 0.24\pi$

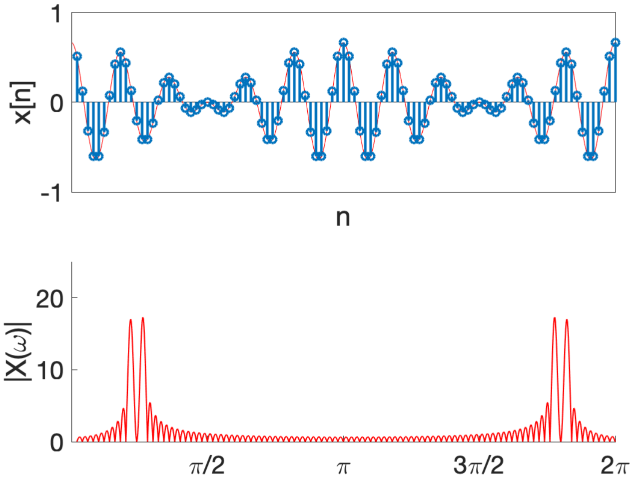

Effect of windowing - spectral smoothing

$x[n] = 1/3(\cos(\omega_1 n) + \cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, with $n = 0, 1, \cdots N-1$, $N = 100$ and $\omega_2 = 0.22\pi$

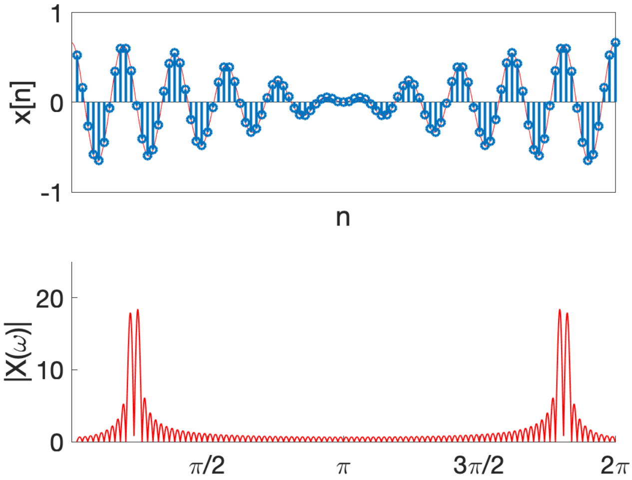

Effect of windowing - spectral smoothing

$x[n] = 1/3(\cos(\omega_1 n) + \cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, with $n = 0, 1, \cdots N-1$, $N = 100$ and $\omega_2 = 0.21\pi$

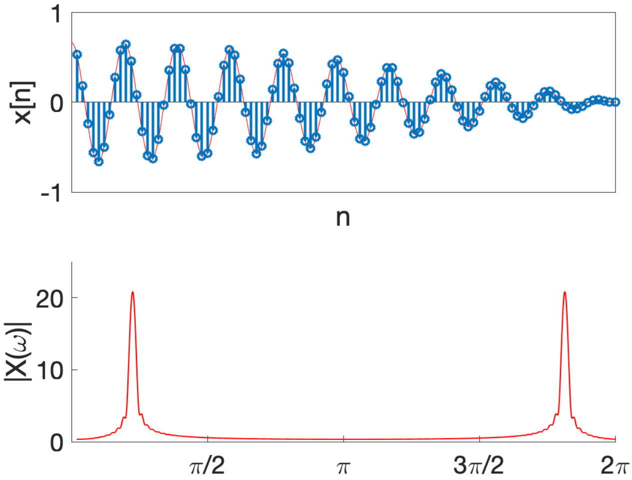

Spectral resolution and the rectangular window

The spectral resolution depends on the width of the main lobe of the window function:

- The spectrum $W(\omega)$ has its first zero-crossing at $\omega = \frac{2\pi}{N}$
- Therefore, two spectral lines ω_1 and ω_2 are not distinguishable if $|\omega_1 - \omega_2| < \frac{2\pi}{N}$.
- If $|\omega_1 - \omega_2| \geq \frac{2\pi}{N}$ we will see two separate lobes in the frequency spectrum.

Thus, the resolution is limited by the number of available samples N. Zero padding will not help.

Effect of windowing - spectral leakage / masking

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots N-1$, $N = 100$ and $A = 1$

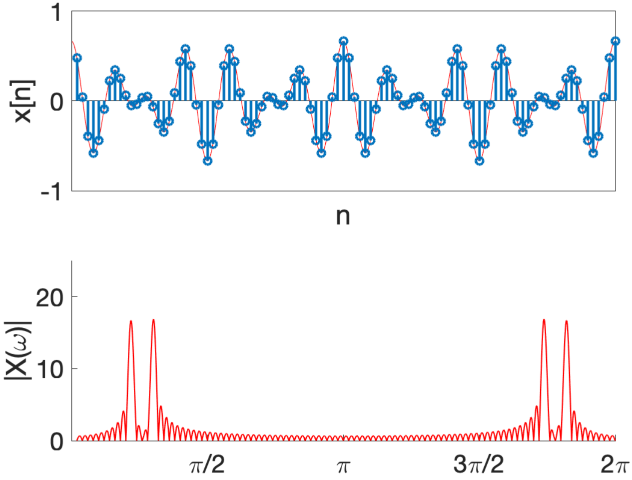

Effect of windowing - spectral leakage / masking

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots N-1$, $N = 100$ and $A = 0.5$

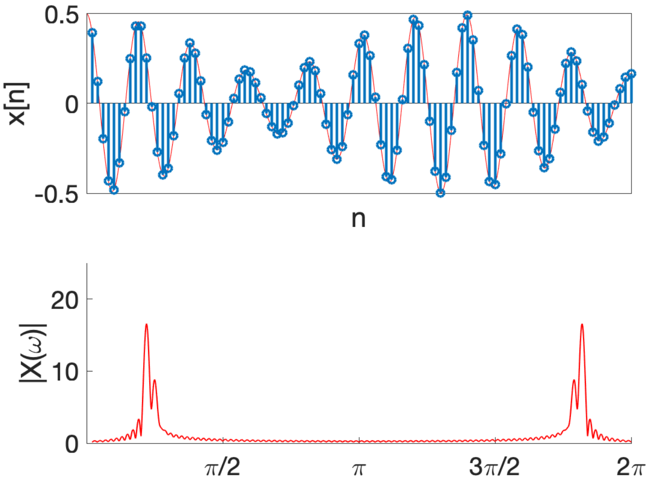

Effect of windowing - spectral leakage / masking

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots N-1$, $N = 100$ and $A = 0.25$

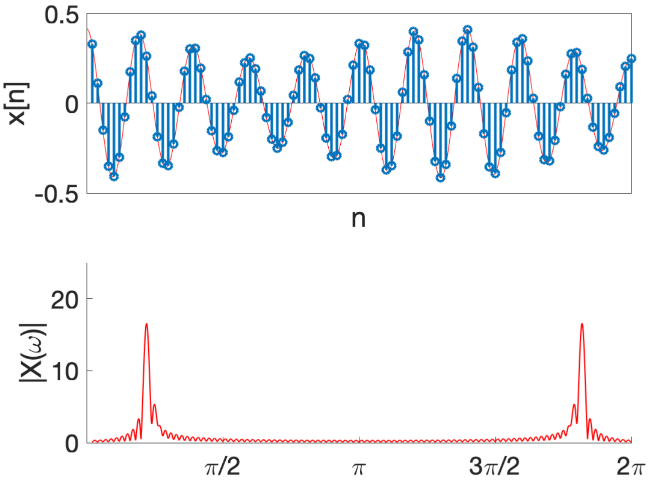

Effect of windowing - spectral leakage / masking

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots N-1$, $N = 100$ and $A = 0.1$

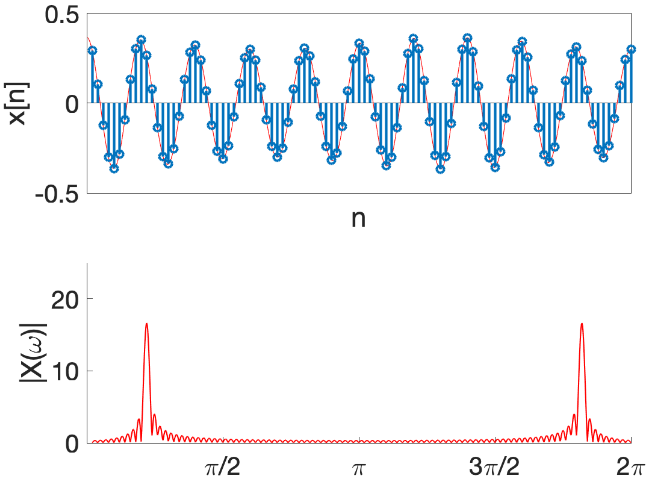

Choice of the window function

We can consider other window functions! Recall from EE2S1 (on FIR filter design):

- Triangular (Bartlett)

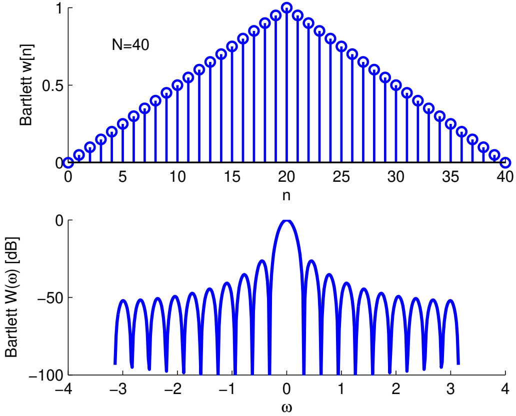

Choice of the window function

We can consider other window functions! Recall from EE2S1 (on FIR filter design):

- Hamming

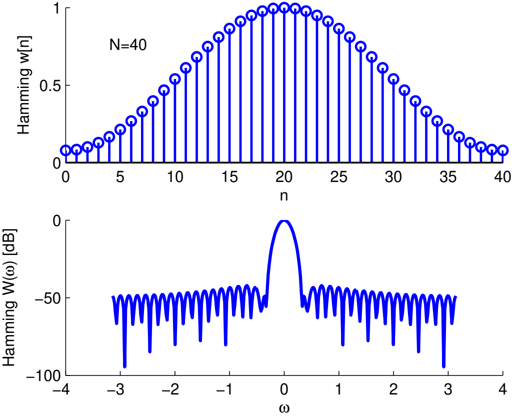

Choice of the window function

We can consider other window functions! Recall from EE2S1 (on FIR filter design):

- Hann

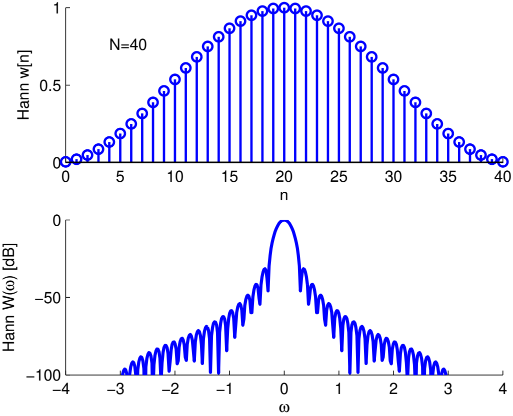

Choice of the window function

We can consider other window functions! Recall from EE2S1 (on FIR filter design):

- Blackman

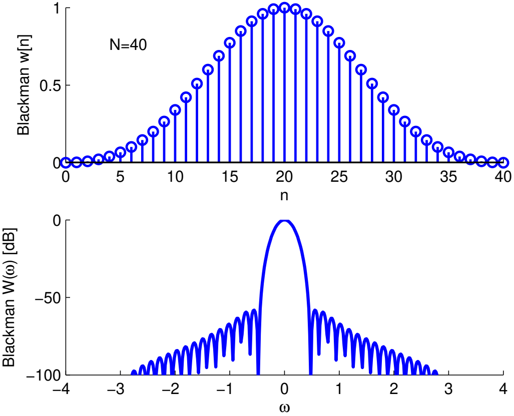

Choice of the window function

We can consider other window functions! Recall from EE2S1 (on FIR filter design):

- Kaiser

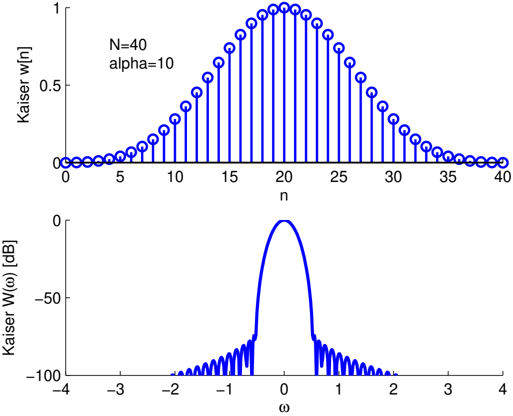

Choice of the window function

In general, there is a trade-off between the width of the main lobe and the amplitude of the sidelobes:

	$\Delta\omega$	ΔA_{MAX} (dB)	sl (dB/oct)
Rectangular	$4\pi/N$	−13	−6
Hamming	$8\pi/N$	−43	−6
Hann	$8\pi/N$	−32	−18
Blackman	$12\pi/N$	−58	−18

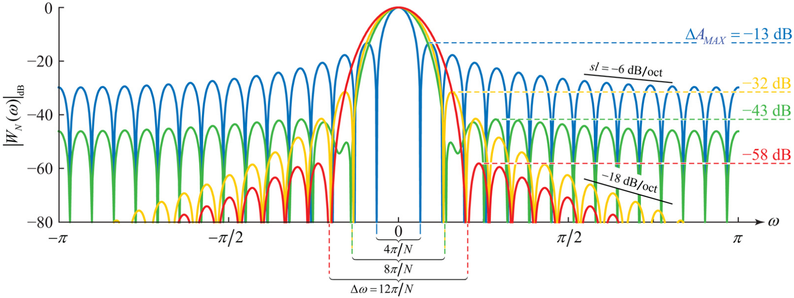

Choice of window - Example 1

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots, N-1$, $N = 100$, $A = 0.1$ and using a rectangular window

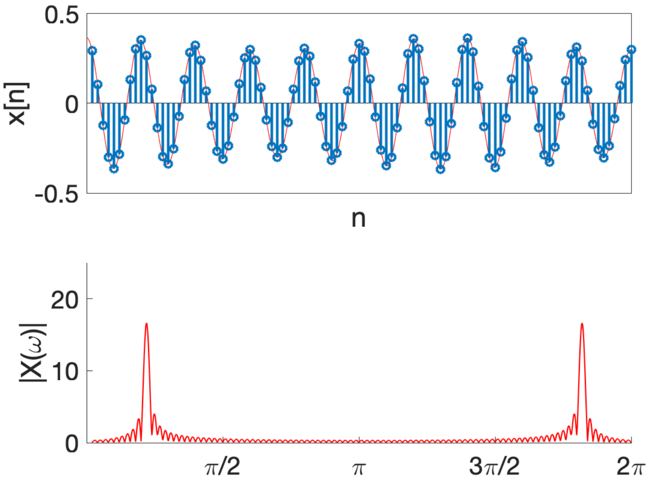

Choice of window - Example 1

$x[n] = 1/3(\cos(\omega_1 n) + A\cos(\omega_2 n))$, $\omega_1 = 0.2\pi$, $\omega_2 = 0.28\pi$ with $n = 0, 1, \cdots, N-1$, $N = 100$, $A = 0.1$ and using a Hann window

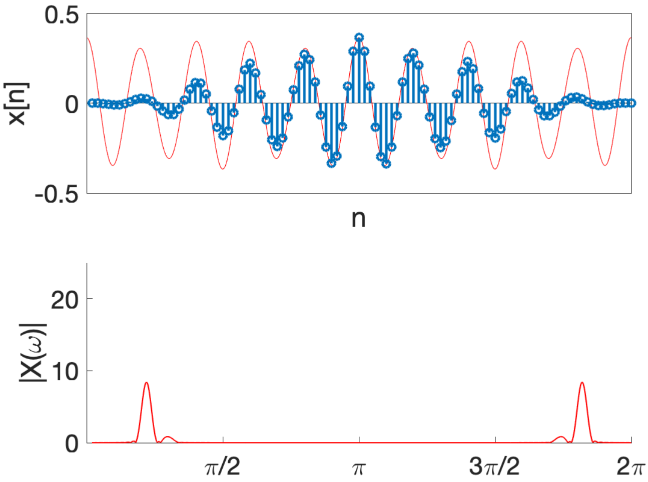

Summary: Spectral analysis using DFT

- In the frequency domain, the spectrum (DTFT) of the finite sequence is equivalent to the convolution of the DTFT of the infinite sequence with the DTFT of the window sequence
- The spectral resolution will depend on the width of the main lobe of the window, which depends on the chosen window function and the number of samples N
- Zero-padding does not increase the spectral resolution but gives a nicer-looking plot

Short-term Fourier Transform (STFT) [Ch. 14.2]

For a non-stationary signal (book: "time-varying signal" (?)), the DFT with a long N masks the changing nature. Remember the EE2S1 *train* and *DTMF* signals; any speech signal.

- Split the signal $x[n]$ into shorter frames (segments) $x_m[n]$ of length L, and apply a window $w_L[n]$:

$$x_m[n] = x[n]w_L[n - mD]$$

The segments can be partially overlapping ($D =$ frame offset).

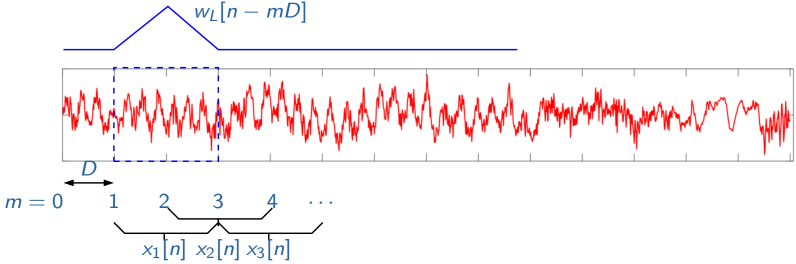

Short-term Fourier Transform (STFT)

- Compute the DFT of each segment

The result is a 2D time-frequency plot $X_L(\omega, m)$: a *spectrogram*

$$X_L(\omega, m) = \sum_{n=-\infty}^{\infty} (x[n]w_L[n - mD])e^{-j\omega n}$$

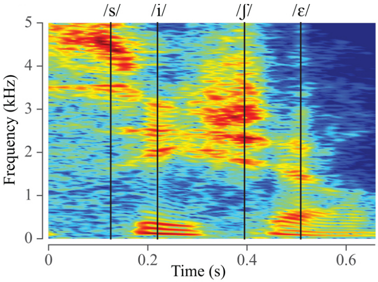

Small L

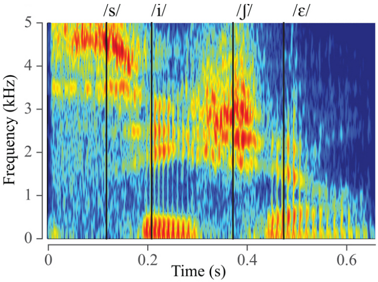

Large L

Constant overlap-add criterion (COLA)

Suppose that the window parameters satisfy $\sum_{m=-\infty}^{\infty} w_L[n - mD] = 1$,

then

$$x[n] = \sum_{m=-\infty}^{\infty} x[n]w_L[n-mD]$$

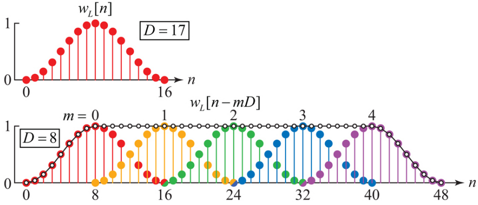

and the STFT is an invertible transform.

$$\sum_{m=-\infty}^{\infty} X_L(\omega, m) = \sum_{m=-\infty}^{\infty} \left(\sum_{n=-\infty}^{\infty} x[n]w_L[n - mD] \right) e^{-j\omega n}$$

$$= \sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n} \quad = X(\omega)$$

Short-term Fourier Transform (STFT)

Many design choices:

- trade-off between the time and frequency resolution:
  large time window $L \Rightarrow$ finer frequency resolution but coarser time resolution
- choice of window function (default is Hann)
- overlap factor $R = \frac{L-1}{D}$, or fractional frame overlap $\frac{R-1}{R}$, e.g., 75%.

Python:

```
stft(x, fs, window='hann', nperseg=256, noverlap, nfft)
```

where $L =$ `nperseg`, $D =$ `nperseg` $-$ `noverlap`, and `nfft` $\geq L$ allows for zero padding.

Power spectral density estimation [Ch. 14.3]

Random signals (stochastic processes) are part of SSP, and the estimation of the power spectral density of a random signal is presented in the final lectures, and analyzed in detail in EE4C03.

The next slides briefly show some connections:

- For a WSS random signal $x[n]$, the DTFT $X(\omega)$ is also random. Therefore, we would look at $\mathrm{E}[|X(\omega)|^2]$.

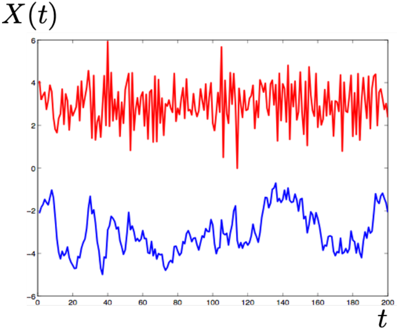

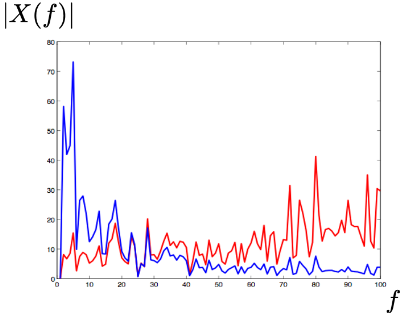

Power spectral density estimation

Since stationary random signals have infinite energy (= do not have a Fourier transform), we have to look at power.

- Let $x_N[n]$ be a window of N samples, then

$$\hat{S}_X(\omega, N) = \frac{1}{N} \left| \sum_{n=0}^{N-1} x_N(n]e^{-j\omega n} \right|^2$$

$$S_X(\omega) = \lim_{N\to\infty} \hat{S}_X(\omega, N)$$

The estimator $\hat{S}_X(\omega, N)$ is called the *periodogram*, and $S_X(\omega)$ is the power spectral density (PSD).

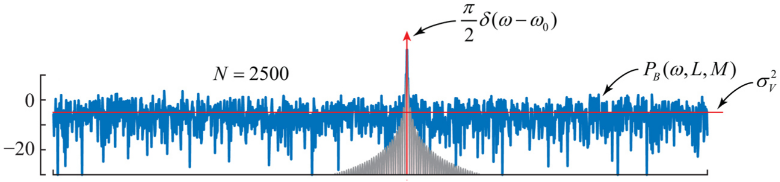

Welch's method

The periodogram is very noisy (has a large variance). To improve, use the STFT!

- Split the signal into M shorter segments of length L (possibly overlapping), and apply a window:

$$x_m[n] = x[n]\, w_L[n - mD]$$

- Compute the DTFT (or rather DFT) of each segment:

$$\hat{S}_{X_m}(\omega, L) = \frac{1}{L} \left| \sum_{n=0}^{L-1} x_m[n] e^{-j\omega n} \right|^2$$

- Average the M segments! This will reduce the variance

$$\hat{S}_X(\omega) = \frac{1}{M} \sum_{m=0}^{M-1} \hat{S}_{X_m}(\omega, L)$$

This is the *averaged modified periodogram*, or the *Welch method*.

Welch's method

Random phase cosine in white noise, using a rectangular window:

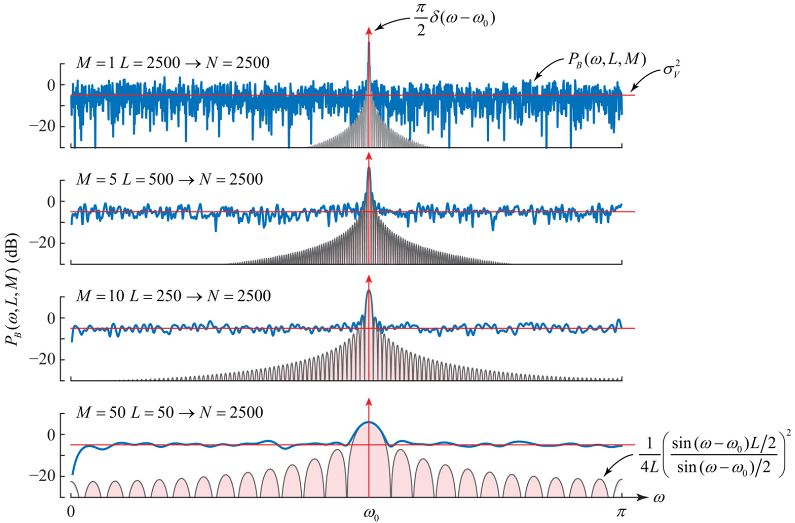

Welch's method

- For a finite number of samples N, which we split into M segments of length L, we have a trade-off between large M (reduces variance) and large L (increases spectral resolution).
- For overlapping segments, data in the segments are not independent, limiting the effect of averaging: keep $D \geq \frac{L}{2}$.
- For proper measure, we must correct for the energy of the window, divide by $\sum_{n=0}^{L-1} w_L^2[n]$.
  For a *density*, also divide by 2π (on an ω-axis) or by F_s (on a frequency axis in Hz).

Python:

```
welch(x, window='hann', nperseg, noverlap, nfft)
```

where $L =$ `nperseg`, $D =$ `nperseg` $-$ `noverlap`, and `nfft` $\geq L$ allows for zero padding.

To do:

- Study the covered parts of chapters 10, 14
- Try to make exercise ...

Next lecture, we look at ADCs.