ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΑΤΡΩΝ
UNIVERSITY OF PATRAS

Signal & Systems

Lecture 9: Windowing

Konstantinos Chatzilygeroudis - costashatz@upatras.gr

Department of Electrical and Computer Engineering
University of Patras

Template made by Panagiotis Papagiannopoulos

Recap: Spectral Leakage

Key idea

- In practice, we observe only a finite-length segment of a signal.
- This is equivalent to multiplying the signal by a window:

$$x_w[n] = x[n]\,w[n]$$

- With a rectangular observation window,

$$w[n] = \begin{cases} 1, & 0 \leqslant n \leqslant N-1 \\ 0, & \text{otherwise} \end{cases}$$

Frequency-domain consequence

- Multiplication in time $\iff$ convolution in frequency
- A pure sinusoid no longer appears as a single spectral line
- Its energy spreads into nearby frequencies:

$$X_w(e^{j\omega}) = X(e^{j\omega}) * W(e^{j\omega})$$

Interpretation

- Sharp truncation at the edges creates spectral spreading
- This spreading is called spectral leakage

Recap: Spectral Leakage (2)

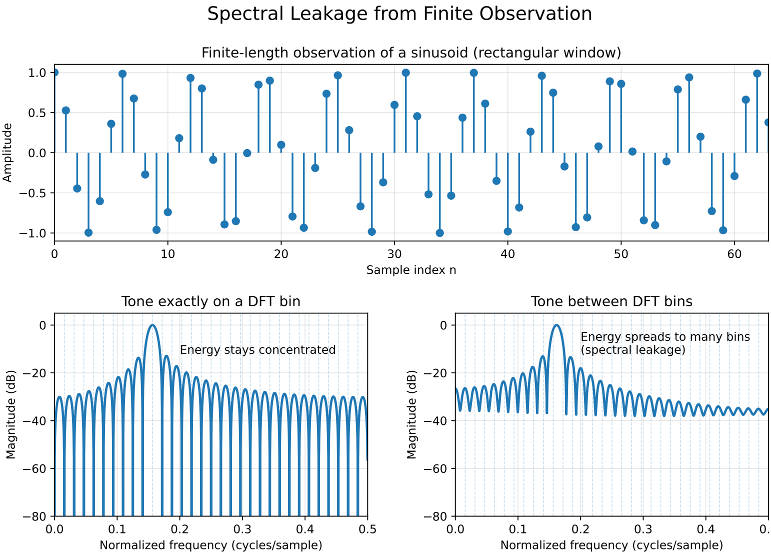

What is Windowing?

Definition

- Windowing means multiplying a signal by a function that selects and shapes a segment:

$$x_w[n] = x[n]\,w[n]$$

Interpretation

- We never observe infinite signals in practice
- We choose how to observe the signal
- The window $w[n]$ defines:
  - what part we keep
  - how smoothly we taper the edges

Window Shapes Example

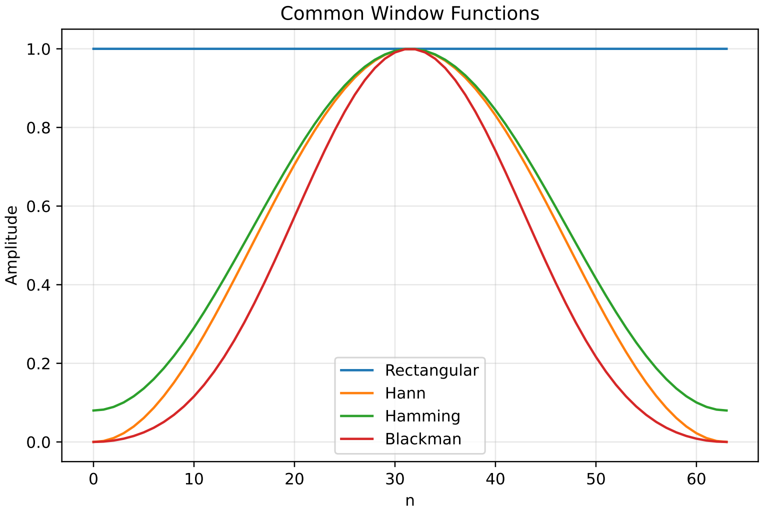

Effect of Windowing in Time

Original signal

- Potentially long or infinite
- May not align with observation boundaries

After windowing

- Signal is truncated and shaped
- Edges become important

Key intuition

- Rectangular window:
  - sharp discontinuities
- Smooth windows:
  - gradual transitions
  - fewer abrupt jumps

Effect of Windowing in Time

Original signal

- Potentially long or infinite
- May not align with observation boundaries

After windowing

- Signal is truncated and shaped
- Edges become important

Key intuition

- Rectangular window:
  - sharp discontinuities
- Smooth windows:
  - gradual transitions
  - fewer abrupt jumps

Why this matters

- Discontinuities → high-frequency components

Effect of Windowing in Time (2)

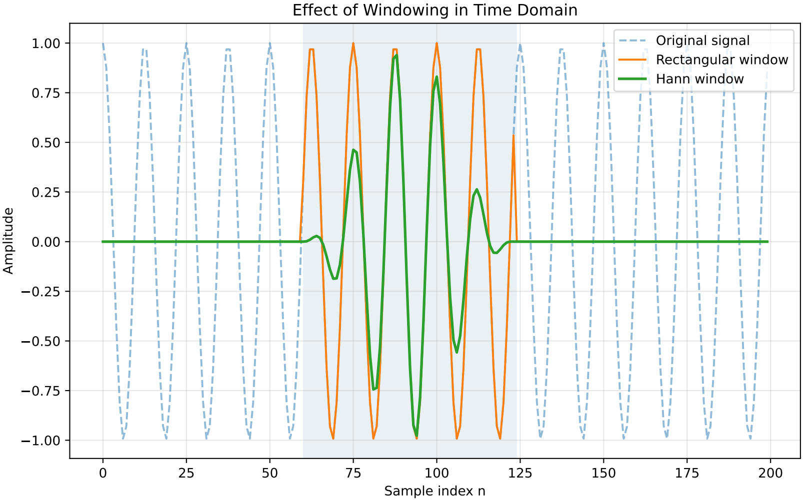

Effect of Windowing in Frequency

Fundamental relation

- Multiplication in time $\iff$ convolution in frequency

$$X_w(e^{j\omega}) = X(e^{j\omega}) * W(e^{j\omega})$$

Interpretation

- The spectrum of the signal is smeared by the window
- The window acts as a spectral filter

Key consequence

- Rectangular window $\rightarrow$ strong side-lobes (more leakage)
- Smooth window $\rightarrow$ lower side-lobes (less leakage)

Big idea

- You do not just observe the signal
- You observe the signal *through the window*

Effect of Windowing in Frequency (2)

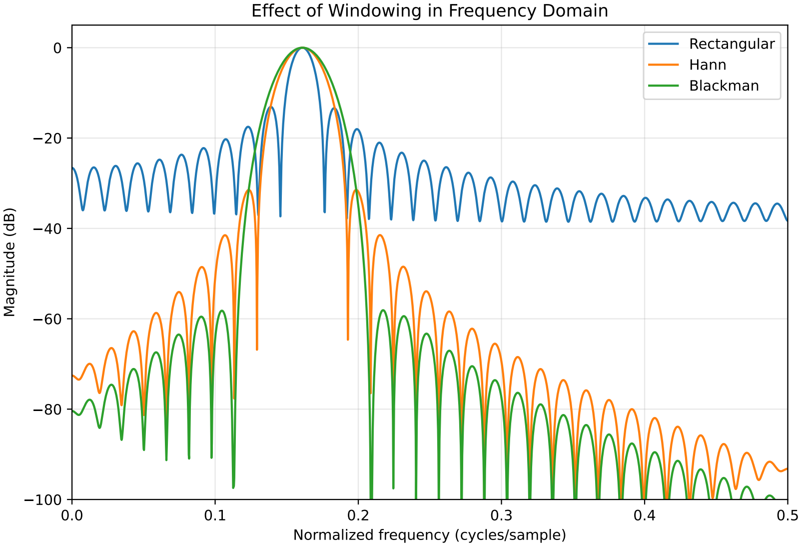

Windowing = Controlled Spectral Tradeoff

By choosing a window, you control:

- Spectral leakage (side-lobes)
- Frequency resolution (main-lobe width)
- Dynamic range (ability to see weak signals)

Key intuition

- Sharp window → narrow main lobe, high side-lobes
- Smooth window → wide main lobe, low side-lobes

Takeaway

Windowing is not just preprocessing — it is a design choice that determines what you can see in the spectrum.

The Fundamental Tradeoff

Choosing a window means balancing two competing goals:

1. **Narrow main lobe** → good frequency resolution
2. **Low side lobes** → low spectral leakage

Visual intuition

- **Main lobe** = how precisely we can locate a frequency
- **Side lobes** = how much energy spills into nearby frequencies

Key point

- We cannot optimize both simultaneously
- A sharper window in time gives:
  - narrower main lobe
  - higher side lobes
- A smoother window in time gives:
  - lower side lobes
  - wider main lobe

Windowing is a controlled tradeoff.

The Fundamental Tradeoff (2)

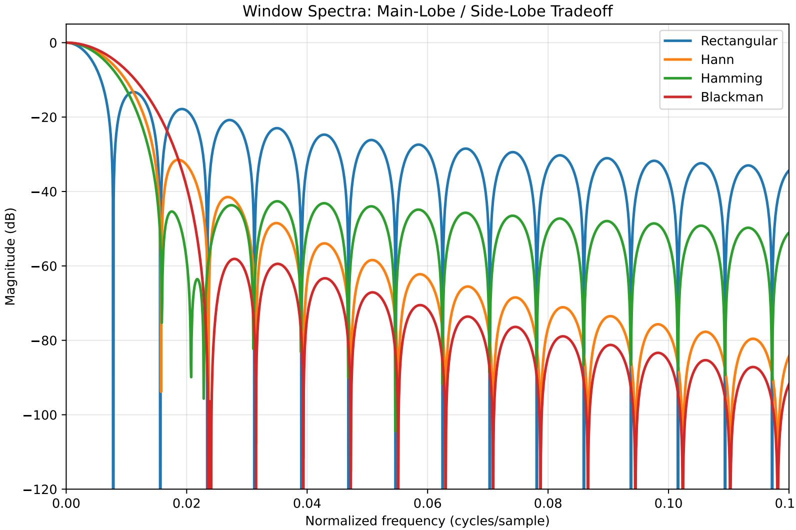

Why Different Windows Behave Differently

Fundamental relation

$$x_w[n] = x[n]\,w[n]$$
$$X_w(e^{j\omega}) = X(e^{j\omega}) * W(e^{j\omega})$$

Interpretation

- Windowing in time means convolution in frequency
- Each window has its own spectral signature
- The observed spectrum is the original spectrum blurred by the window spectrum

Consequences

- Narrow spectral kernel → better resolution
- Low side-lobe kernel → less leakage
- Different windows behave differently because $W(e^{j\omega})$ is different

Takeaway

You do not observe the signal alone; you observe the signal *through the spectrum of the window*.

Why Different Windows Behave Differently (2)

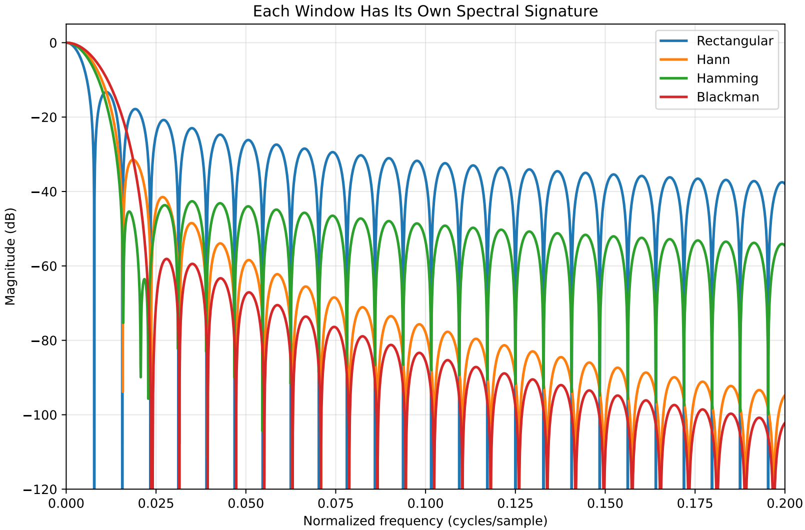

Common Window Definitions I

For $0 \leqslant n \leqslant N-1$, the windowed signal is

$$x_w[n] = x[n]\, w[n].$$

Rectangular window

$$w_{\mathrm{rect}}[n] = 1$$

Hann window

$$w_{\mathrm{Hann}}[n] = \frac{1}{2}\left(1 - \cos\left(\frac{2\pi n}{N-1}\right)\right)$$

Hamming window

$$w_{\mathrm{Hamming}}[n] = 0.54 - 0.46\cos\left(\frac{2\pi n}{N-1}\right)$$

Observation

Rectangular keeps all samples equally. Hann and Hamming taper the edges smoothly toward zero or near-zero.

Common Window Definitions II

For $0 \leqslant n \leqslant N-1$:
Blackman window

$$w_{\text{Blackman}}[n] = 0.42 - 0.5\cos\left(\frac{2\pi n}{N-1}\right) + 0.08\cos\left(\frac{4\pi n}{N-1}\right)$$

Kaiser window

$$w_{\text{Kaiser}}[n] = \frac{I_0\left(\beta\sqrt{1-\left(\frac{2n}{N-1}-1\right)^2}\right)}{I_0(\beta)}$$

where $I_0(\cdot)$ is the modified Bessel function of the first kind.

- $\beta = 0$ gives a rectangular-like shape
- Larger $\beta \rightarrow$ smoother taper, lower side-lobes, wider main lobe

Why Kaiser is useful

It gives a tunable tradeoff between resolution and leakage.

Many Windows Are Cosine Sums

Many practical windows can be written as

$$w[n] = \sum_{m=0}^{M} a_m \cos\left(\frac{2\pi mn}{N-1}\right), \qquad 0 \leqslant n \leqslant N-1.$$

Examples:

- Rectangular: only a constant term
- Hann: constant + one cosine term
- Hamming: constant + one cosine term with different coefficients
- Blackman: constant + two cosine terms

Interpretation

More shaping freedom in time → more control over side-lobes and main-lobe width in frequency.

Comparing Common Windows

Window	Main Lobe Width	Side Lobes	Dynamic Range	Typical Use
Rectangular	Narrowest	Worst leakage	Poor	Frequency estimation for clean, well-aligned tones
Hann	Wider	Much lower	Good	General-purpose spectral analysis
Hamming	Similar to Hann	Slightly lower first side-lobe	Good–better	Audio / practical FFT analysis
Blackman	Wide	Very low	High	Weak tones near strong interferers
Kaiser	Adjustable	Adjustable	Adjustable	Tunable engineering compromise

Key takeaway

Windows move you along the resolution–leakage–dynamic-range tradeoff curve.

Windowing and Usable Dynamic Range

Punchline

Your window determines your usable dynamic range.

Why?

- A strong sinusoid produces side-lobes
- These side-lobes act like an artificial floor
- A weak sinusoid must rise above that floor to be visible

Interpretation

- High side-lobes → poor dynamic range
- Low side-lobes → high dynamic range

Visual intuition

- **Main lobe** tells us how well we separate nearby tones
- **Side lobes** tell us whether weak tones get buried

Important

- Better dynamic range usually comes with a wider main lobe
- So improving detectability often reduces resolution

Windowing and Usable Dynamic Range (2)

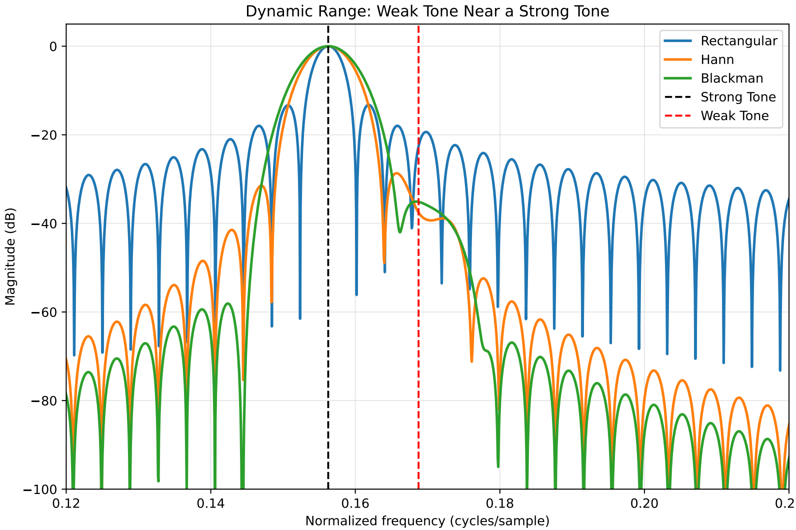

Side-Lobes as an Artificial Floor

Key idea

- A strong sinusoid does not remain perfectly localized in the spectrum
- Because of the window, it creates a main lobe and side lobes
- The side lobes form an artificial floor around the tone

Interpretation

- If a weak tone lies below that floor, it becomes hard to detect
- Therefore, side-lobe level directly affects usable dynamic range

Question to ask

Can I see a signal that is 40 dB below a stronger one?

Rule of thumb

- Higher side lobes $\rightarrow$ poorer dynamic range
- Lower side lobes $\rightarrow$ better weak-signal visibility

Side-Lobes as an Artificial Floor (2)

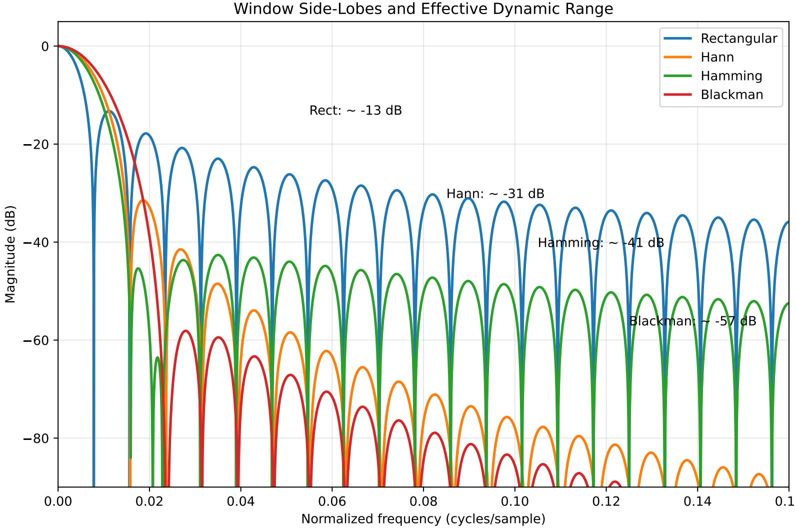

Approximate Side-Lobe Floors of Common Windows

Window	Main Lobe	First Side-Lobe Level	Dynamic Range Intuition
Rectangular	Narrow	≈ -13 dB	Poor
Hann	Medium	≈ -31 dB	Good
Hamming	Medium	≈ -41 dB	Better
Blackman	Wide	≈ -57 dB	Excellent

Interpretation

A window with a side-lobe level near -40 dB is much more suitable for revealing tones that are 30 to 40 dB below a strong one.

Windows with lower side lobes provide higher usable dynamic range.

How Do We Compute the Side-Lobe Level?

Let $w[n]$ be a window of length N, and let

$$W[k] = \sum_{n=0}^{N-1} w[n] e^{-j2\pi kn/N_{\mathrm{FFT}}}$$

be a densely sampled DFT of the window.

Step 1: Normalize by the peak

$$\widetilde{W}[k] = \frac{|W[k]|}{\max_m |W[m]|}$$

Step 2: Convert to dB

$$W_{\mathrm{dB}}[k] = 20 \log_{10}\big(\widetilde{W}[k]\big)$$

Step 3: Find the largest side-lobe peak

- Exclude the main-lobe region around $k = 0$
- Search for the highest local maximum outside the main lobe

Definition

$$\text{Side-lobe level} = \max_{k \in \text{side-lobes}} W_{\mathrm{dB}}[k]$$

This value is usually reported relative to the main-lobe peak, in dB.

Worked Example: Hann Window

For the Hann window,

$$w[n] = \frac{1}{2}\left(1 - \cos\left(\frac{2\pi n}{N-1}\right)\right), \qquad 0 \leqslant n \leqslant N-1$$

Procedure

1. Compute a high-resolution FFT of $w[n]$
2. Normalize so the peak is at $0\ \mathrm{dB}$
3. Ignore the main lobe near DC
4. Measure the highest remaining peak

Result

$$\text{First side-lobe level} \approx -31\ \mathrm{dB}$$

Meaning

A strong tone observed with a Hann window creates leakage about $31\ \mathrm{dB}$ below its peak in the first side-lobe.

Worked Example: Hann Window (2)

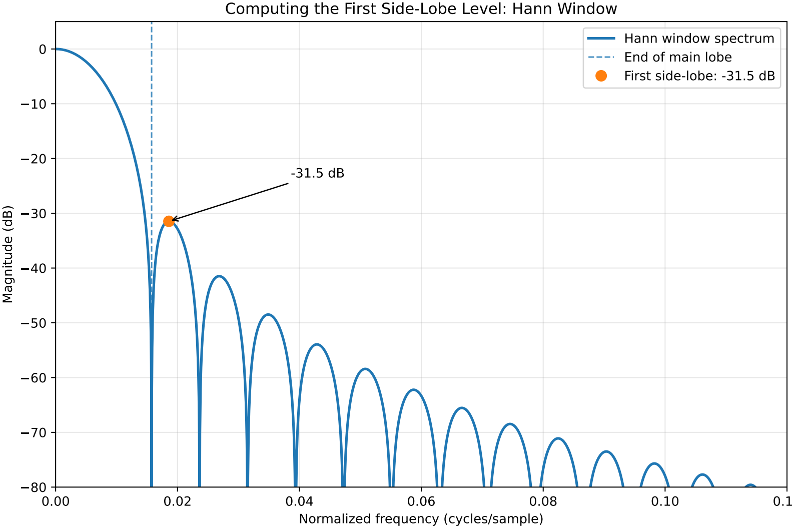

Analytical Example: Side-Lobe of the Rectangular Window

For the rectangular window

$$w[n] = 1, \qquad 0 \leqslant n \leqslant N-1,$$

its DTFT is

$$W(e^{j\omega}) = \sum_{n=0}^{N-1} e^{-j\omega n} = e^{-j\omega(N-1)/2}\,\frac{\sin(N\omega/2)}{\sin(\omega/2)}.$$

Ignoring the linear phase term, the normalized magnitude is

$$\frac{|W(e^{j\omega})|}{N} = \left|\frac{\sin(N\omega/2)}{N\sin(\omega/2)}\right|.$$

For large N, define $x = \frac{N\omega}{2}$, so that near the main lobe / first side-lobe region,

$$\frac{|W(e^{j\omega})|}{N} \approx \left|\frac{\sin x}{x}\right|.$$

The side-lobe peaks occur where

$$\frac{d}{dx}\left(\frac{\sin x}{x}\right) = 0 \implies \frac{x\cos x - \sin x}{x^2} = 0 \implies x\cos x - \sin x = 0 \implies \tan x = x.$$

The first solution after $x = \pi$ is

$$x_1 \approx 4.493.$$

Therefore, the first side-lobe amplitude is $\left|\frac{\sin x_1}{x_1}\right| \approx 0.217$, and its level relative to the main-lobe peak is

$$20\log_{10}(0.217) \approx -13.26\ \mathrm{dB}.$$

Summary: What a Window Really Does

- Windowing changes the spectrum of the observation
- Every window imposes a different balance between:
  - **Resolution** (main-lobe width)
  - **Leakage suppression** (side-lobe level)
  - **Dynamic range** (visibility of weak tones)
- No single window is best for all tasks

Rule of thumb

- Need sharp separation of close tones? → narrower main lobe
- Need to reveal weak tones near strong ones? → lower side lobes

Windowing Changes Amplitude

Observation

- When we apply a window, measured amplitudes decrease
- The same sinusoid gives different peak values for different windows

Why does this happen?

- Windowing multiplies the signal:

$$x_w[n] = x[n]\,w[n]$$

- Most windows have values less than 1
- So the signal energy is attenuated

Key idea

Windowing does not just reshape the spectrum — it also scales the amplitude.

Question

- Can we correct for this amplitude loss?

Coherent Gain

Definition

The coherent gain (CG) of a window is

$$\mathrm{CG} = \frac{1}{N}\sum_{n=0}^{N-1} w[n].$$

Interpretation

- It is the average value of the window
- It tells us how much a coherent sinusoid is scaled

Special case

- Rectangular window:

$$\mathrm{CG} = 1$$

Meaning

A sinusoid aligned with a DFT bin is scaled approximately by the coherent gain.

Amplitude Correction Using Coherent Gain

Measured vs true amplitude

- Let A_{true} be the true amplitude
- Let A_{measured} be what we observe after windowing

Relation

$$A_{\text{measured}} \approx \mathrm{CG} \cdot A_{\text{true}}$$

Correction

$$A_{\text{true}} \approx \frac{A_{\text{measured}}}{\mathrm{CG}}$$

Practical rule

Always divide by the coherent gain when estimating amplitudes from a windowed FFT.

In dB

$$A_{\text{dB, corrected}} = A_{\text{dB}} - 20 \log_{10}(\mathrm{CG})$$

Typical Coherent Gain Values

Window	Coherent Gain
Rectangular	1.00
Hann	0.50
Hamming	0.54
Blackman	0.42

Interpretation

- Hann reduces amplitude by about 50%
- Blackman reduces amplitude even more

Tradeoff reminder

- Better leakage suppression → lower coherent gain

Takeaway

Windows that improve dynamic range also attenuate amplitude more.

Coherent Gain Example

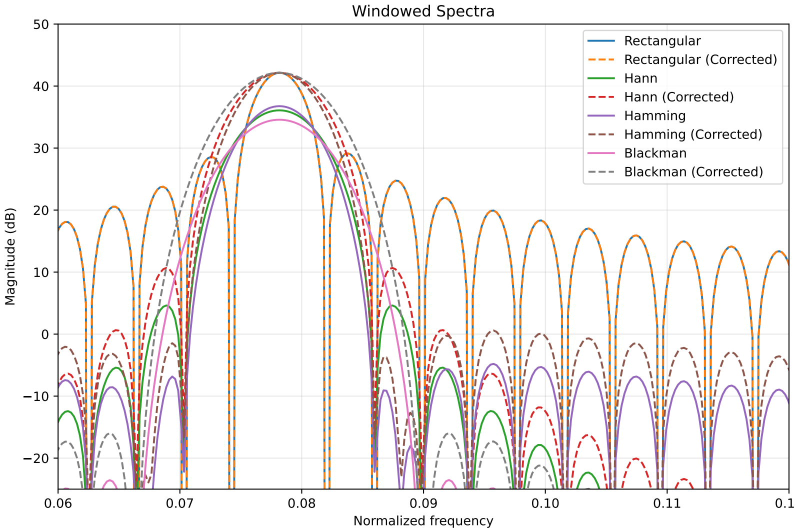

Experiment: Same Signal, Different Windows

Goal

- Compare how different windows affect what we can observe in the spectrum

Signal

$$x[n] = A_1 \cos(2\pi f_1 n) + A_2 \cos(2\pi f_2 n), \qquad A_2 \ll A_1$$

with:

- a **strong** tone at frequency f_1
- a nearby **weak** tone at frequency f_2

Question

- Can we see the weak tone in the presence of the strong one?

What we will compare

Rectangular vs Hann vs Blackman

What Different Windows Reveal

Rectangular window

- Narrow main lobe
- High side-lobes
- Leakage from the strong tone can hide the weak tone

Hann window

- Lower side-lobes
- Weak tone becomes visible
- Main lobe is wider

Blackman window

- Very low side-lobes
- Weak tone is very clear
- Peaks are more blurred due to wider main lobe

Key lesson

Window choice determines what you can see.

What Different Windows Reveal (2)

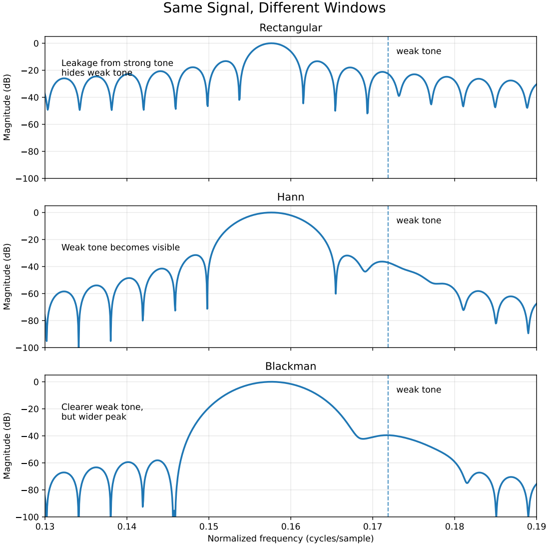

Experiment: Close Frequencies

Goal

- Study how window choice affects the ability to separate nearby tones

Signal

$$x[n] = \cos(2\pi f_1 n) + \cos(2\pi f_2 n), \qquad f_1 \approx f_2$$

with:

- two **equal-amplitude** sinusoids
- very **small frequency separation**

Question

- Can the spectrum resolve the two tones as distinct peaks?

What we compare

Rectangular window vs Hann window

Experiment: Resolution Tradeoff

Rectangular window

- Narrower main lobe
- Better frequency resolution
- Two close tones can still appear as separate peaks
- But the spectrum is more affected by leakage / ripple

Hann window

- Lower side-lobes
- Smoother spectrum
- Wider main lobe
- Very close tones may merge into one peak

Key lesson

Reducing leakage usually comes at the cost of frequency resolution.

Experiment: Resolution Tradeoff (2)

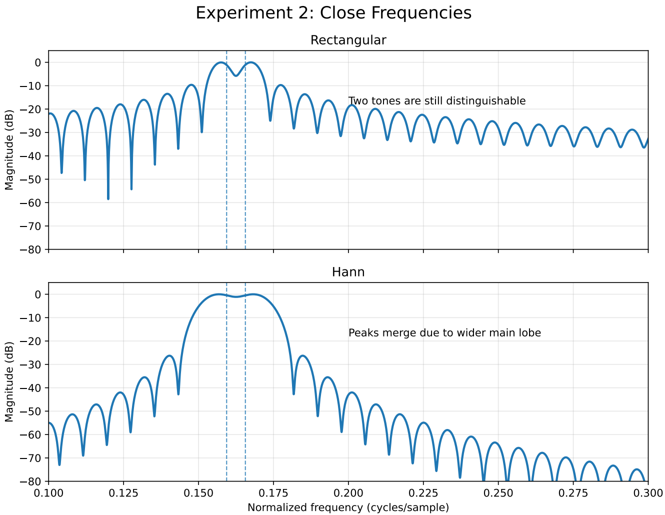

Experiment 3: Two-Tone Dynamic Range Test

Goal

- Test how well different windows reveal a weak tone near a strong one

Signal

$$x[n] = A_1 \cos(2\pi f_1 n) + A_2 \cos(2\pi f_2 n)$$

with:

- a strong tone at $0\ \mathrm{dB}$
- a weak nearby tone at:

$$-20\ \mathrm{dB}, \quad -40\ \mathrm{dB}, \quad -60\ \mathrm{dB}$$

Question

- As the weak tone gets smaller, which windows still reveal it?

What we compare

Rectangular vs Hann vs Blackman

Prediction Before We Look

Recall

- Rectangular window has narrow main lobe but high side-lobes
- Hann has lower side-lobes
- Blackman has very low side-lobes

Expected behavior

- **Rectangular:** weak tone disappears quickly
- **Hann:** visible up to roughly 30–40 dB difference
- **Blackman:** visible even near 60 dB difference

What determines visibility?

The weak tone must rise above the side-lobe floor created by the strong tone.

Same Geometry, Different Dynamic Range

What happens as the weak tone decreases?
Rectangular

- High side-lobes from the strong tone
- Weak tone is quickly buried

Hann

- Lower leakage floor
- Weak tone remains visible longer

Blackman

- Very low side-lobes
- Weak tone can remain visible even at very low levels

Key lesson

Dynamic range is not only a property of the signal – it is also a property of the window.

Visualization

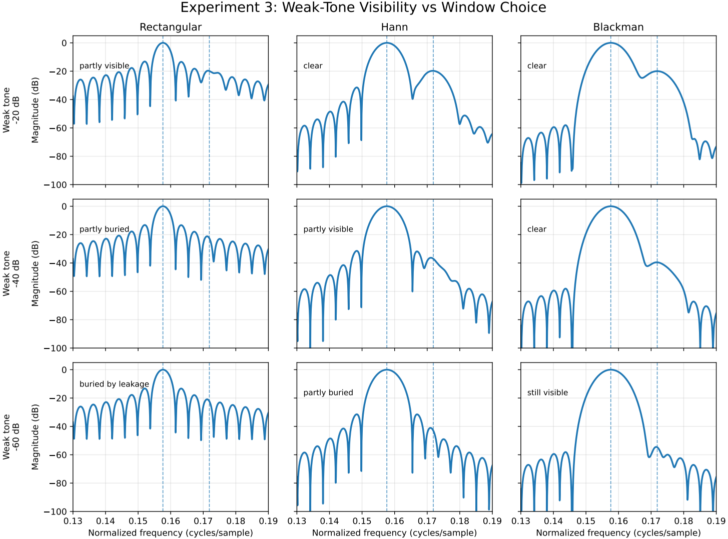

Key Question

Did the signal disappear

– or did our tool fail to see it?

Did the signal disappear

– or did our tool fail to see it?

- The signal is still there in all cases
- What changes is the analysis window
- Different windows create different side-lobe floors
- That floor determines whether weak components remain visible

Takeaway

A spectrum analyzer does not simply reveal truth; it reveals truth through a chosen window.

Thank you

- **Any Questions?**
- **Office Hours**:
  - **Mon & Tue** *(09:00-11:00)*
  - 24/7 by email (`costashatz@upatras.gr`, subject: *ECE_SS_AM*)
- **Material and Announcements**

Laboratory of Automation & Robotics