Digital Signal Processing

Curated undergraduate study track (not accredited, no external exam). Original explanations and practice only, following the reviewed course handout.

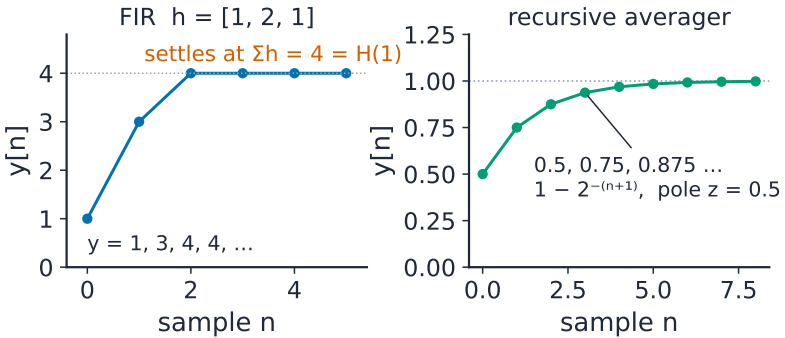
Refounding signal theory in samples

- Foundations first (systems, **DTFT**/**DFT**, **FFT**), then filters (FIR windows, IIR bilinear), then rates and a verified project.
- Every design measured against its specification; every speedup an operation count, not folklore.

Discrete systems: one list decides everything

1. **FIR** h = [1, 2, 1] on a unit step: y = 1, 3, 4, 4 … settling at $\Sigma h = 4$, the DC gain H(1); a pulse [1,1,1] reads 1, 3, 4, 3, 1 - the taps mirrored.
2. **IIR** The averager y[n] = (y[n−1] + x[n])/2 steps up 0.5, 0.75, 0.875, 0.9375 … = $1 - 2^{-(n+1)}$, because its pole sits at z = 0.5 - it approaches, never arrives.
3. **Stability** BIBO asks one question of the taps: does $\Sigma|h[n]|$ converge? For rational systems that is exactly all poles inside |z| = 1.

The sliding table and the pole

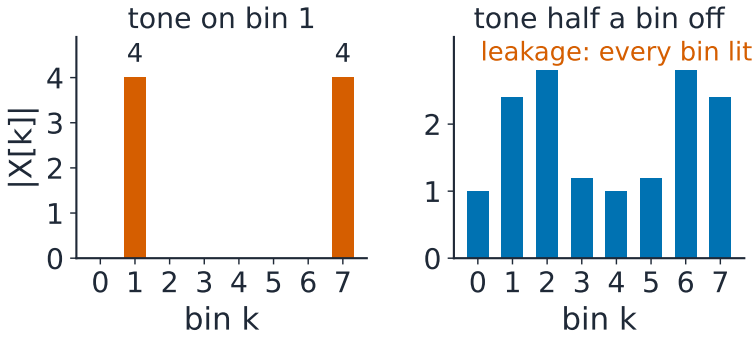


The FIR settles in three samples at $\Sigma h = 4$; the recursive averager closes on 1 as $1 - 2^{-(n+1)}$ from its $z = 0.5$ pole.

DTFT and DFT: sampling in frequency

1. **On-bin** 8-point **DFT** of a cosine on bin 1: $|X| = 4$ at bins 1 and 7, zero elsewhere - the periodic copies join seamlessly.
2. **Off-bin** Half a bin off: 2.40, 2.80, 1.19, 1.0 ... in every bin - the rectangle sidelobes at -13 dB paint the spectrum; Hamming drops them to -43 dB.
3. **Padding** Zero-padding interpolates the same spectrum - it resolves nothing. **Resolution** was decided when you chose how long to measure.

Two spikes become every bin

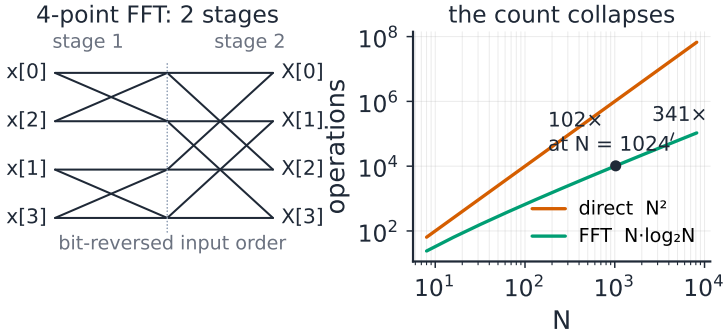


On bin 1: two clean spikes of 4. Half a bin off: every bin lit - **leakage** is the rectangular window speaking.

The FFT: same answer, fewer operations

1. **Count** $N = 1024$: direct $N^2 = 1\,048\,576$ operations, **FFT** $N \cdot \log_2 N = 10\,240$ - $102\times$ fewer; $N = 4096$: $341\times$, and the gap widens with N.
2. **Real time** 1 Hz bins at $f_s = 1$ kHz: $N = 1024$ gives 0.977 Hz; at 10 ns per operation the transform takes ~100 µs against a 1.024 s record.
3. **Map** Bin 100 of that **DFT** is $100 \times 0.977 = 97.66$ Hz - the bin-to-hertz map $k \cdot f_s/N$, the classic lab error when forgotten.

The butterfly ladder

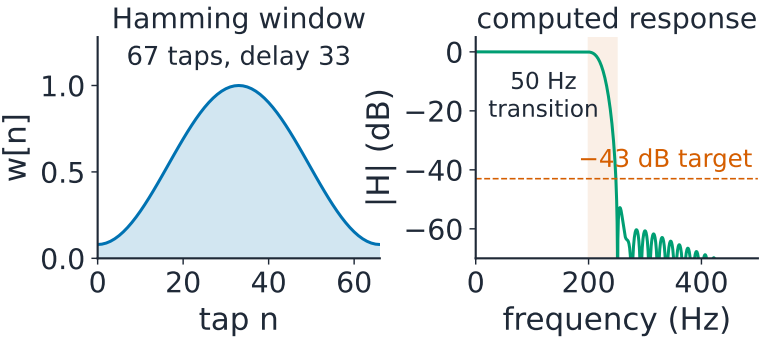


Two stages of butterflies build a 4-point **FFT**; at N = 1024 the count collapses 102× and keeps falling.

FIR: the window is the fate

1. **Size** Hamming at $f_s = 1$ kHz, 50 Hz transition: $N = 3.3 \cdot 1000/50 = 66$ - take 67 (odd, Type-I), delay 33 samples; $N = 31$ gives 106 Hz: these are transition-width estimates.
2. **Verify** Measure the response from the taps: target stopband at or below -43 dB and ripple inside ± 0.05 dB. A -37 dB result fails: inspect taps, cutoff, window, N and measurement band.
3. **Optimal** Windows are convenient, not optimal - the alternation theorem behind Parks-McClellan makes equiripple the optimum.

Window in, response out

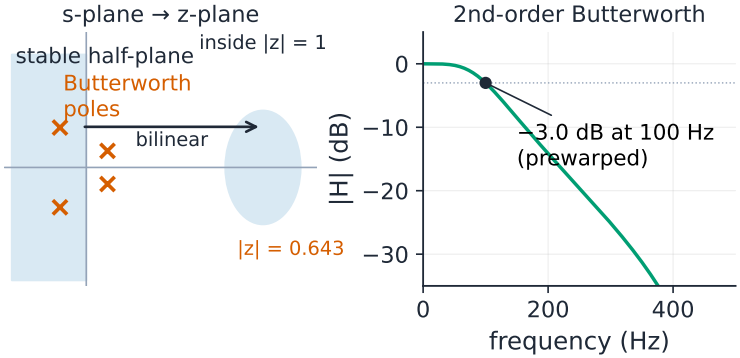


67 Hamming taps, 225 Hz ideal cutoff: computed stopband from 250 Hz peaks near −52.3 dB; delay 33 samples.

IIR: analog lore, one substitution

1. **Biquad** 2nd-order Butterworth at 100 Hz, $f_s = 1$ kHz: **prewarp** $\omega_a = 2000 \cdot \tan(0.1\pi) = 649.8$ rad/s, $K = 3.078$, $H(z) = (1+2z^{-1}+z^{-2})/(14.82 - 16.94z^{-1} + 6.12z^{-2})$.
2. **Verify** DC gain $4/4 = 1.000$; $|H|$ at 100 Hz $= -3.0$ dB - the cutoff lands exactly; poles at $|z| = 0.643$, inside by the mapping itself.
3. **Trap** Skip the prewarp and the cut lands at 96.9 Hz; set $T = 1$ too ($K = 1$) and it cuts at 250 Hz, reading a complacent -0.05 dB at 100 Hz - silently off-spec.

The left half-plane moves in



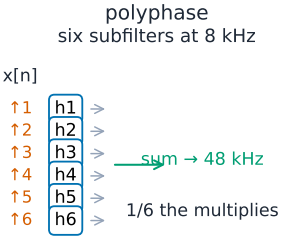
The stable half-plane folds inside $|z| = 1$ (poles land at 0.643) and the **prewarp** puts the -3 dB point on 100 Hz exactly.

Multirate: work where it is cheap

1. **Up** 8 kHz speech to 48 kHz: interpolate by 6, low-pass at 4 kHz kills images at 8k, 16k, 40k; **polyphase** runs 6 subfilters at 8 kHz - 1/6 the multiplies.
2. **Ratio** 48 $\rightarrow$ 72 kHz is L $= 3$, M $= 2$: upsample to 144 kHz, ONE low-pass at 24 kHz stops the $\uparrow 3$ images and the $\downarrow 2$ aliases, keep every second sample.
3. **Gain** Zero-stuffing loses a factor L of amplitude - compensate it in the filter; forgetting the gain of L is the classic multirate bug.

Images every 8 kHz





Stuffing 5 zeros paints images every 8 kHz; one low-pass at 4 kHz erases them; **polyphase** pays 1/6 the multiplies.

The project: specify, verify, measure

1. **Spec** 1 Hz bins at $f_s = 1$ kHz needs $N \geq 1000$ - already 1.0 s against the <1 s latency line: the conflict surfaces at design time, exactly when you want it.
2. **Stages** Filter $\rightarrow$ window $\rightarrow$ **FFT** $\rightarrow$ threshold, each verified alone: a defect found in a stage costs an hour; in the integrated system, a week of bisection.
3. **Report** The final table lists spec, measured, margin per line - 5 of 6 and honest beats 6 of 6 claimed with 4 measured.

Four blocks and a margin table

verify each stage alone — then measure the whole



spec	measured	margin
1 Hz	0.977 Hz	ok
<1 s	1.024 s	✗
−10 dB	−10 + 27 dB	ok

Each block carries its own check; the pipeline ends in the spec-versus-measured margin table with no silent goals.

Check yourself

- FIR: h = [1,2,1] on a unit step settles at? (4 - the tap sum, the DC gain.)
- **DFT**: 1 Hz bins at $f_s = 1$ kHz need N = ? (1000 - one second of data.)
- **FFT**: N = 1024 speedup over direct? (102× - N^2 against $N \cdot \log_2 N$.)
- **IIR**: design for 100 Hz at 1 kHz with T = 1 (K = 1) and it cuts at? (250 Hz - the classic silent off-spec; skipping only the prewarp gives 96.9 Hz.)