FFT Peak Interpolation: Quadratic vs. Quartic Fits, and the Zero-at-Center / Worst-near-Quarter-Bin Bias Pattern

1 The problem

A tone sampled by an N-point FFT almost never lands exactly on a bin center. Writing the true frequency as $k_0 + \delta$ bins, with $\delta \in [-0.5, 0.5]$ the fractional offset from the nearest bin k_0, the three complex DFT values around the peak are

$$X_{-1} = X(k_0 - 1), \qquad X_0 = X(k_0), \qquad X_{+1} = X(k_0 + 1). \tag{1}$$

The interpolation problem is: given only these three complex numbers, estimate δ as accurately as possible.

2 Approach 1: fit the complex spectrum directly (quartic)

X_{-1}, X_0, X_{+1} are complex, so one natural idea is to fit a quadratic to the real parts $\operatorname{Re} X(k)$ and a second, independent quadratic to the imaginary parts $\operatorname{Im} X(k)$, each as a function of bin index k. With exactly three points, each fit is an exact (Lagrange) interpolation, not a least-squares fit:

$$R(x) = a_R x^2 + b_R x + c_R, \qquad R(-1) = \operatorname{Re} X_{-1},\ R(0) = \operatorname{Re} X_0,\ R(1) = \operatorname{Re} X_{+1}, \tag{2}$$

$$I(x) = a_I x^2 + b_I x + c_I, \qquad I(-1) = \operatorname{Im} X_{-1},\ I(0) = \operatorname{Im} X_0,\ I(1) = \operatorname{Im} X_{+1}. \tag{3}$$

The power at any (non-integer) offset x is then

$$P(x) = R(x)^2 + I(x)^2, \tag{4}$$

a product of two quadratics squared and summed – a genuine **4th-order polynomial (quartic) in** x. Finding its true maximum means solving $P'(x) = 0$, a cubic equation, not reading off a single vertex formula the way a quadratic allows.

This is a legitimate way to pose the problem, but it has a structural weakness: a 3-point exact fit has no information about how the spectrum behaves *beyond* those three points. Nothing forces $R(x)$ or $I(x)$ to decay away from the peak, so the quartic $P(x)$ is not guaranteed to be well-behaved (single-peaked) over the interval – it can have its true maximum land on the wrong side of the measured peak entirely as δ moves away from zero. Figure 1 shows this numerically: bias exceeding half a bin by $\delta = 0.5$, far worse than simply not correcting at all.

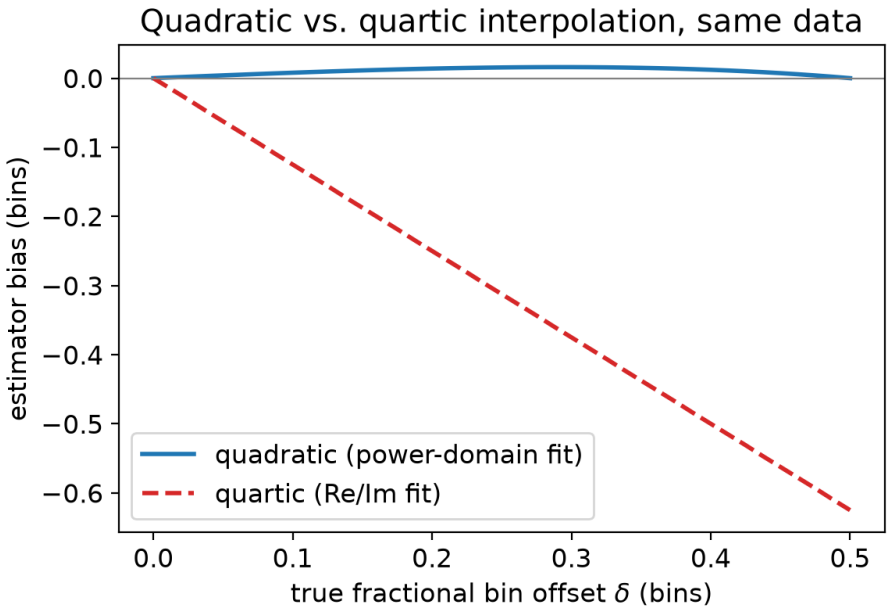


Figure 1: The power-domain quadratic fit (Eq. (7)) compared against the quartic Re/Im fit from this section, on identical data. The quartic approach is not a refinement of the quadratic one – it is substantially worse, because the unconstrained 3-point complex fit has no mechanism to stay well-behaved away from the measured bins.

3 Approach 2: fit the power directly (quadratic)

The practical alternative sidesteps the complex-plane fit entirely and interpolates the (always real, always non-negative) **power**

$$y_1 = |X_{-1}|^2, \qquad y_2 = |X_0|^2, \qquad y_3 = |X_{+1}|^2 \tag{5}$$

with a single quadratic. Because $|X(k)|^2$ for a windowed single tone has a smooth, unimodal mainlobe shape with no sign changes, a quadratic is a far more honest local model of it than it is of the oscillating, phase-dependent $\mathrm{Re}\,X(k)$ or $\mathrm{Im}\,X(k)$ individually. The standard 3-point vertex formula gives the fractional correction directly, no cubic to solve:

$$\delta \;=\; \frac{1}{2}\,\frac{y_1 - y_3}{y_1 - 2y_2 + y_3}. \tag{6}$$

A linear-power fit like Eq. (6) still assumes a parabolic mainlobe shape that no real window (rectangular or Hann) actually has. Fitting the quadratic to the *logarithm* of the power instead matches a windowed mainlobe's real shape far better:

$$\delta \;=\; \frac{1}{2}\,\frac{\ln y_1 - \ln y_3}{\ln y_1 - 2\ln y_2 + \ln y_3}. \tag{7}$$

This is also a reason the power-domain fit is better-posed than the complex-plane one in the first place: $|X(k)|^2$ is (to good approximation) **independent of the tone's arbitrary phase** at the start of the capture window, while $\mathrm{Re}\,X(k)$ and $\mathrm{Im}\,X(k)$ individually are not. Two captures of the

same tone at the same frequency but different starting phase give identical y_1, y_2, y_3, but different (R, I) triples – so a fit built directly on R and I is modeling something that depends on a quantity (δ and the phase) the quadratic has no way to separate.

4 Why the bias is zero at the bin center

Equation (7) makes one fact immediate: if the true frequency sits exactly on a bin ($\delta = 0$), the spectrum is symmetric about k_0, so $y_1 = y_3$ exactly. The numerator $\ln y_1 - \ln y_3$ is then exactly zero, giving $\delta_{\text{est}} = 0$ with **no approximation involved** – this holds regardless of noise, window choice, or N, purely from the symmetry of the three points being fit.

5 Why the bias returns to zero at the bin boundary, and peaks near a quarter bin

The same symmetry argument extends to $\delta = 0.5$ (a tone sitting exactly between two bins): for a symmetric window, the spectrum is symmetric about the midpoint between k_0 and $k_0 + 1$, so the 3-point log-power fit again returns zero bias at that exact point. Between these two zeros, the bias is not zero – it rises from 0, peaks, and falls back to 0, and from symmetry considerations alone the extremum is expected near the midpoint of the interval, i.e. around $\delta \approx 0.25$. Figure 2 confirms this numerically for the shipped Hann-window, log-power estimator at $N = 2048$: zero bias at $\delta = 0$ and $\delta = 0.5$, and a peak bias of about 0.016 bins near $\delta \approx 0.29$ – in close agreement with real bench measurements of the same formula (-0.078 Hz at the mid-bin point, corresponding to the same $\sim$0.016-bin scale).

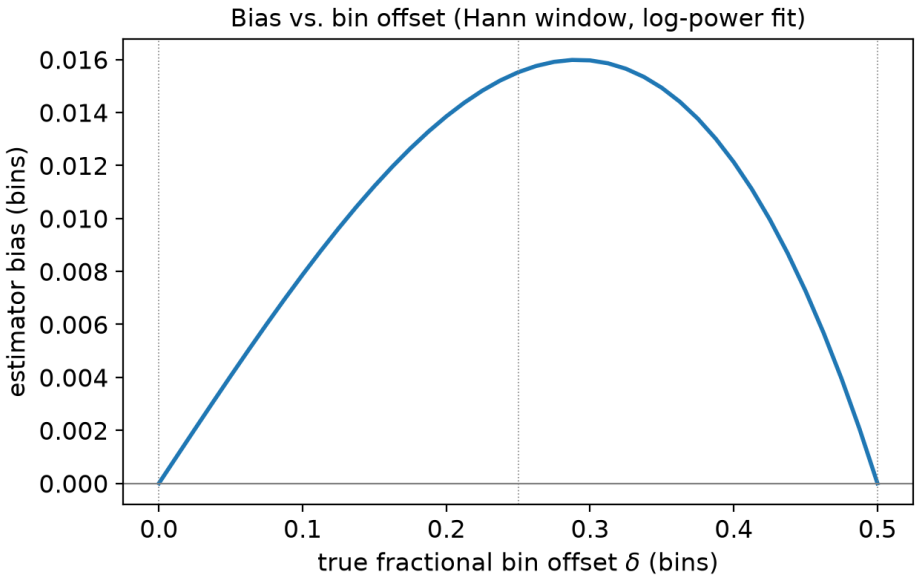


Figure 2: Bias of the Hann-windowed, log-power 3-point estimator as a function of the true fractional bin offset δ. Bias is exactly zero at the bin center ($\delta = 0$) and at the bin boundary ($\delta = 0.5$), with the worst case near a quarter bin.

6 Comparing estimators

Table 1 puts numbers behind the discussion above, evaluating every estimator this investigation considered on the same noiseless single-tone model used to generate Figures 1 and 2 ($N = 2048$, $F_s = 10\,\text{kHz}$, bin width $4.883\,\text{Hz}$, peak swept across $\delta \in [0, 0.5]$ in steps of 0.025 bins). "Peak bias" is the worst-case error over that sweep; "bias at $\delta = 0.25$" isolates the quarter-bin point used elsewhere in this document as the representative worst-case location for the shipped estimator.

Table 1: Interpolation bias by estimator, noiseless single tone, $N = 2048$. Rows are grouped by window; the log-power quadratic (Eq. (7)) is the one shipped in firmware.

Estimator	Window	Peak bias (bin)	Peak bias (Hz)	Bias @ δ=0.25
Linear-power quadratic (Eq. (6))	rect	0.284	1.385	−0.231
Log-power quadratic (Eq. (7))	rect	0.167	0.817	−0.156
Jacobsen	rect	0.000	0.000	0.000
Jacobsen + Candan 2011 correction	rect	0.000	0.000	0.000
Jacobsen + Candan 2013 correction	rect	0.000	0.000	0.000
Log-power quadratic (SHIPPED)	Hann	**0.016**	**0.078**	**0.016**
Jacobsen	Hann	0.250	1.222	−0.125
Jacobsen + Candan 2011 correction	Hann	0.250	1.222	−0.125
Quinn's 2nd estimator	Hann	0.667	3.257	−0.500
Quartic Re/Im fit (rejected)	Hann	0.625	3.052	−0.313

Three things stand out:

- **Jacobsen and its Candan corrections are exact for a rectangular window** on this noiseless single-tone model – their closed-form derivations assume exactly that window, and the table confirms zero bias to numerical precision. They are *not* derived for a Hann window, and simply carrying the rectangular-window Candan correction factor over to Hann-windowed bins (as tested here) does not fix anything: the peak bias stays at 0.25 bin, identical to uncorrected Jacobsen, because the correction constant itself is wrong for this window.
- **The shipped Hann + log-power quadratic has the smallest peak bias of every Hann-window estimator tested** – over 15× smaller than Jacobsen on the same window, and more than 40× smaller than the rejected quartic fit. This is the quantitative version of the symmetry argument in the sections above: because $\ln|X(k)|^2$ is (to a good approximation) a genuine parabola in k for a Hann mainlobe, the log-power quadratic is close to exact precisely where the others are not derived to be exact at all.
- **Quinn's 2nd estimator and the quartic Re/Im fit are the two worst performers on a Hann window** – worse than not correcting at all for Quinn's estimator (0.667 bin exceeds the 0.5-bin range being corrected for), and worse than half a bin for the quartic fit. Both share the same root cause described in Section 2: a 3-point fit with no constraint forcing it to decay away from the peak, applied to a window whose shape the estimator's derivation did not assume.

7 Summary

- Fitting quadratics to $\mathrm{Re}\, X(k)$ and $\mathrm{Im}\, X(k)$ separately and combining them produces a quartic power function whose true maximum requires solving a cubic – and in practice is badly behaved away from the measured bins.
- Fitting a single quadratic (or log-quadratic) directly to the power $|X(k)|^2$ is both simpler and better motivated: power is phase-invariant and has a smooth, unimodal shape that a quadratic genuinely approximates well locally.
- The log-power fit's bias is exactly zero at the bin center and at the bin boundary, by symmetry of the three points being fit, with the worst case in between, near a quarter bin.
- Table 1 confirms this holds up against the standard literature estimators, not just the rejected quartic fit: on a Hann window, the shipped log-power quadratic has the smallest peak bias of every estimator tested, since Jacobsen, Candan, and Quinn's closed-form estimators are all derived for a rectangular window and are not valid as-is on Hann-windowed bins.

References

- E. Jacobsen, P. Kootsookos, "Fast, Accurate Frequency Estimators," *IEEE Signal Processing Magazine*, 2007.
- B. G. Quinn, "Estimating Frequency by Interpolation Using Fourier Coefficients," *IEEE Trans. Signal Processing*, 1994.
- Ç. Candan, "A Method For Fine Resolution Frequency Estimation From Three DFT Samples," *IEEE Signal Processing Letters*, vol. 18, no. 6, 2011 (and the refined arctan correction, *IEEE Signal Processing Letters*, vol. 20, no. 9, 2013).
- F. J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform," *Proc. IEEE*, 1978.