Fast Equal Division Phase Window Fourier Analysis

Mathematical Foundation, Fast Chirp-Convolution Realization, Comparative Spectral Analysis and Computational Trade-offs

Technical White Paper

8 October 2026

Purpose. This document presents the modern mathematical formulation of the Equal Division Phase Window (EDPW) method, originally conceived in 1997 for high-rate spectral analysis. It distinguishes the original transform from its fast implementation, derives its exact relation to the conventional DFT/FFT, establishes an exact transform-operator distance, develops an algebraically equivalent chirp-convolution realization, and evaluates both full-spectrum and selective computational regimes. The analysis adds the Direct-EDPW and Fast-EDPW crossover analysis, separates principal convolution cost from an optimized total runtime estimate, introduces a unified N–w complexity map, and explicitly couples computational economy to the spectral information required by the task. The comparative campaign separates arithmetic economy, spectral fidelity, and detection sensitivity as distinct engineering properties. This development phase adds comparative foundations for exact and approximate spectral algorithms and a reproducibility-oriented spectral-vulnerability benchmark.

Research draft. Claims of computational advantage are restricted to the conditions explicitly analyzed herein.

Contents

1 Motivation and Historical Engineering Problem **2**

2 Conventional DFT and FFT Reference **3**

3 Original EDPW Definition **3**

4 Exact Relationship Between EDPW and the DFT **4**

4.1 Exact convergence at $w = N$ 5

5 Transform-Matrix Form and Exact Structural Distance **5**

6 Direct EDPW Arithmetic Cost and Crossover **6**

7 Fast EDPW by Exact Chirp-Convolution Factorization **7**

7.1 Equivalence theorem 7

8 Fast EDPW Complexity and Crossover **8**

8.1 Principal count versus optimized total estimate 8

8.2 Direct versus Fast realization 9

9 Unified Complexity–Spectral Trade-off **9**

10 Selective Frequency-Region Processing **10**

11 Full-Spectrum and Selective Operating Regions **11**

12 Additional Theoretical Foundations **12**

12.1 Frequency-grid theorem 12

12.2 Single-tone response and finite-window selectivity 12

12.3 Detection is not spectral equality 13

12.4 Selective-complexity crossover theorem 13

12.5 Exact fast/direct equivalence and error decomposition 14

12.6 Initial AWGN and coherent-processing implication 14

12.7 Open theoretical question 14
12.8 Computational frequency-grid limit 15

13 Numerical Experiments **15**

13.1 Experiment 0: introductory pure-sinusoid frequency localization 15
13.2 Window weighting as a dispersion-control trade-off 16
13.3 Experiment I: corrected selective wideband workload 17
13.4 Experiment II: computational frequency-grid scalability at fixed machine capacity 18

14 Comparative Spectral Algorithms: Mathematical Background **21**

14.1 Full DFT and radix-2 FFT 21
14.2 Direct EDPW and exact Fast-EDPW 22
14.3 Output-pruned FFT 22
14.4 Goertzel recurrence 22
14.5 Chirp-Z transform (CZT) 22
14.6 Short-time Fourier transform (STFT) 23

15 Expanded Total Computational Complexity Comparison **23**

16 Experiment III: Comparative Spectral Vulnerability **24**

16.1 Common-bin comparison across the spectral methods 24
16.2 STFT: a separate time-localization experiment 27
16.3 Additional evidence: spectral fidelity, arithmetic cost, and blind discovery . . 28

17 Proposed Distance Metrics for the Numerical Campaign **30**

17.1 Complex spectral distance 30
17.2 Magnitude-spectrum distance 31
17.3 Maximum normalized bin error 31
17.4 Task-level metrics 31

18 Experimental Status and Remaining Engineering Work **32**

19 Limitations and Claims Deliberately Not Made **32**

20 Conclusions **32**

21 References **34**

A Reference Equations **35**

B Public Code Coverage and Verification Protocol **35**

B.1 Cold-call versus reusable-kernel arithmetic 35

C Reproducibility Status **36**

Abstract

The Equal Division Phase Window (EDPW) method partitions an N-sample record into equal windows of length w while retaining the global N-point Fourier phase reference. Each time window is paired with a corresponding set of global frequency indices and produces independent coefficients. The direct full-spectrum formulation requires Nw complex kernel products, while an exact chirp-convolution factorization reduces the principal convolution work to $O(N \log w)$. A detailed crossover analysis shows, however, that asymptotic notation alone does not determine the preferred realization: for sufficiently small w, Direct EDPW can require fewer principal complex multiplications than a full radix-2 FFT, and can also be cheaper than the chirp-convolution realization. Fast-EDPW can enter a full-spectrum arithmetic-win region at sufficiently large N, but its phase/chirp products must be distinguished from the principal convolution count. These computational gains do not imply FFT-like spectral reconstruction. The coefficient grid remains F_s/N, whereas finite-window selectivity is governed approximately by F_s/w. The appropriate engineering design is therefore a constrained optimization: choose w and the Direct/Fast realization to minimize arithmetic subject to a required detection, depiction, or reconstruction quality. Selective processing remains a major opportunity, while comparative experiments must establish whether lower arithmetic preserves the spectral information required by each task. The resulting spectral dispersion is the explicit price of short-window economy and motivates a scan–detect–zoom architecture in which EDPW first identifies candidate regions and detailed analysis is then concentrated only where required.

1. Motivation and Historical Engineering Problem

Consider spectral analysis over a very wide instantaneous bandwidth. For a real-valued signal spanning approximately 0 to 2 GHz, a sampling rate of at least 4 GS/s is required by the Nyquist criterion; practical front-end margin can move the rate toward 5 GS/s. At 5 GS/s, a 16-bit stream carries 80 Gbit/s (10 GB/s), while a 32-bit stream carries 160 Gbit/s (20 GB/s). The difficulty is therefore not only the arithmetic count of a Fourier transform. Acquisition, buffering, memory bandwidth, data movement, latency, and the fraction of the frequency range that actually requires detailed examination all become first-order design variables.

The EDPW concept was originally motivated by this type of problem. Its defining idea is to preserve a global Fourier frequency grid associated with a long record N, but to calculate a coefficient using only a corresponding local group of w samples. This produces independent time-frequency blocks. Such a structure is not automatically superior to a full FFT when every output is required; its potential advantage appears when only selected frequency regions must be inspected, for example during threshold-based spectral search followed by local refinement.

This paper deliberately separates three questions:

1. What exactly is the EDPW transform and how does it differ mathematically from the DFT?
2. Can the original EDPW coefficients be calculated faster without changing their definition?
3. Under what workload does the resulting architecture offer a useful computational advantage?

2. Conventional DFT and FFT Reference

For a record $x[n]$, $n = 0, \ldots, N-1$, the unnormalized N-point DFT is

$$X_F[m] = \sum_{n=0}^{N-1} x[n] \mathrm{e}^{-j2\pi nm/N}, \qquad m = 0, \ldots, N-1. \tag{1}$$

The associated frequency grid is

$$f_m = \frac{mF_s}{N}, \qquad \Delta f = \frac{F_s}{N}. \tag{2}$$

For N a power of two, a common radix-2 count for complex butterfly multiplications is [3]

$$C_{\mathrm{FFT}} = \frac{N}{2} \log_2 N. \tag{3}$$

This count is used as a transparent reference in the present paper; implementation-specific real-operation counts can be introduced in later hardware studies.

It is essential to distinguish grid spacing from physical resolving capability. Equation (2) states where the DFT samples the frequency axis. Resolution of nearby components depends on observation duration, windowing, signal model, noise, and the detection criterion.

3. Original EDPW Definition

Assume

$$N = Kw, \tag{4}$$

where w is the EDPW window length and $K = N/w$ is the number of windows. Define

$$I_k = \{kw, kw+1, \ldots, kw+w-1\}, \qquad k = 0, \ldots, K-1. \tag{5}$$

For a frequency index $m \in I_k$, the EDPW coefficient is defined as

$$X_E[m] = \frac{N}{w} \sum_{n \in I_k} x[n] \mathrm{e}^{-j2\pi nm/N}. \tag{6}$$

The factor N/w is an amplitude compensation reflecting that only w samples contribute to a coefficient instead of all N samples.

Three properties define EDPW:

1. The kernel retains the *global* denominator N, not w.
2. Window k produces only the frequency indices $m \in I_k$.
3. Outputs from different windows are independent; no summation of window spectra is performed.

Consequently, EDPW is not equivalent to applying an ordinary w-point DFT independently to each segment.

4. Exact Relationship Between EDPW and the DFT

For $m \in I_k$, partition the conventional DFT coefficient into contributions inside and outside the corresponding EDPW window:

$$X_F[m] = S_k[m] + R_k[m], \tag{7}$$

$$S_k[m] = \sum_{n \in I_k} x[n] \mathrm{e}^{-j2\pi nm/N}, \tag{8}$$

$$R_k[m] = \sum_{n \notin I_k} x[n] \mathrm{e}^{-j2\pi nm/N}. \tag{9}$$

From Eq. (6),

$$X_E[m] = \frac{N}{w} S_k[m]. \tag{10}$$

Therefore the exact coefficient error is

$$\boxed{E_w[m] = X_E[m] - X_F[m] = \left(\frac{N}{w} - 1\right) S_k[m] - R_k[m].} \tag{11}$$

No statistical assumption is used in Eq. (11). It identifies the two sources of departure from the DFT: scaling of the retained local contribution and omission of all samples outside the paired window.

A direct triangle-inequality bound follows:

$$|E_w[m]| \le \left|\frac{N}{w} - 1\right| \sum_{n\in I_k} |x[n]| + \sum_{n\notin I_k} |x[n]|. \tag{12}$$

This bound is conservative because it discards phase cancellation, but it is valid for arbitrary input.

4.1. Exact convergence at $w = N$

If $w = N$, there is one window, $I_0 = \{0, \ldots, N-1\}$. Then $N/w = 1$ and $R_0[m] = 0$ for every m. Hence

$$\boxed{X_E[m] = X_F[m] \quad \text{for all } m \text{ when } w = N.} \tag{13}$$

Thus the DFT is an exact limiting member of the EDPW family.

5. Transform-Matrix Form and Exact Structural Distance

Let $F_N \in \mathbb{C}^{N\times N}$ denote the unnormalized DFT matrix,

$$[F_N]_{m,n} = \mathrm{e}^{-j2\pi mn/N}. \tag{14}$$

Define $A_w \in \mathbb{C}^{N\times N}$ by

$$[A_w]_{m,n} = \begin{cases} \dfrac{N}{w}\mathrm{e}^{-j2\pi mn/N}, & \lfloor m/w \rfloor = \lfloor n/w \rfloor, \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

Then

$$\boldsymbol{X}_F = F_N \boldsymbol{x}, \qquad \boldsymbol{X}_E = A_w \boldsymbol{x}, \tag{16}$$

and

$$\boldsymbol{X}_E - \boldsymbol{X}_F = (A_w - F_N)\boldsymbol{x}. \tag{17}$$

The matrix contains Nw entries inside the K diagonal blocks and $N^2 - Nw$ entries outside. Since each DFT entry has unit magnitude,

$$\|A_w - F_N\|_F^2 = Nw\left(\frac{N}{w} - 1\right)^2 + (N^2 - Nw) \tag{18}$$

$$= N^2\left(\frac{N-w}{w}\right). \tag{19}$$

Therefore

$$\boxed{\|A_w - F_N\|_F = N\sqrt{\frac{N-w}{w}}.} \tag{20}$$

Because $\|F_N\|_F = N$, the normalized structural distance is

$$\boxed{D_{\mathrm{op},F}(w) = \frac{\|A_w - F_N\|_F}{\|F_N\|_F} = \sqrt{\frac{N}{w} - 1.}} \tag{21}$$

Equation (21) is an *exact Frobenius distance between transform matrices*. It is not a worst-case relative output-error formula and should not be interpreted as a direct prediction of carrier-frequency error, magnitude-spectrum error, or detection probability for a particular signal.

For a signal-dependent norm bound, Eq. (17) gives

$$\|\boldsymbol{X}_E - \boldsymbol{X}_F\|_2 \le \|A_w - F_N\|_2 \, \|\boldsymbol{x}\|_2 \le \|A_w - F_N\|_F \, \|\boldsymbol{x}\|_2. \tag{22}$$

The Frobenius bound is safe but generally loose.

6. Direct EDPW Arithmetic Cost and Crossover

Each of the $K = N/w$ blocks calculates w outputs from w samples. Direct full-spectrum evaluation therefore requires

$$C_{\mathrm{direct}} = \frac{N}{w} w^2 = Nw. \tag{23}$$

Using throughout this paper the same radix-2 principal complex-multiplication reference

$$C_{\mathrm{FFT}} = \frac{N}{2} \log_2 N, \tag{24}$$

Direct EDPW is arithmetically cheaper than the full FFT whenever

$$Nw < \frac{N}{2} \log_2 N \quad \Longleftrightarrow \quad \boxed{w < \frac{1}{2} \log_2 N.} \tag{25}$$

This corrects an over-general conclusion that can result from examining only moderate and large windows. For example, at $N = 1024$, the FFT reference is 5,120 principal complex multiplications: Direct EDPW requires 2,048 at $w = 2$ and 4,096 at $w = 4$, but 8,192 at $w = 8$. Thus a genuine full-spectrum Direct-EDPW arithmetic-win region exists for sufficiently small w.

The crossover is computational, not spectral. A w-sample EDPW coefficient has a first-null scale approximately F_s/w; in units of the global FFT-bin spacing F_s/N, this is approximately N/w bins. Under Eq. (25), $N/w > 2N/\log_2 N$, so the economically favorable Direct-EDPW

region is intrinsically coarse relative to an N-point FFT. Consequently, Direct-EDPW economy must not be presented as generic FFT-quality spectral reconstruction.

7. Fast EDPW by Exact Chirp-Convolution Factorization

The objective is now to accelerate Eq. (6) without altering its coefficients. For block k, set

$$a = kw, \qquad n = a + r, \qquad m = a + s, \tag{26}$$

with $r, s = 0, \ldots, w-1$. The kernel becomes

$$\mathrm{e}^{-j2\pi nm/N} = \mathrm{e}^{-j2\pi(a+r)(a+s)/N} \tag{27}$$

$$= \mathrm{e}^{-j2\pi a^2/N}\mathrm{e}^{-j2\pi ar/N}\mathrm{e}^{-j2\pi as/N}\mathrm{e}^{-j2\pi rs/N}. \tag{28}$$

Use the identity

$$rs = \frac{r^2 + s^2 - (r-s)^2}{2}. \tag{29}$$

Then

$$\boxed{\mathrm{e}^{-j2\pi rs/N} = \mathrm{e}^{-j\pi r^2/N}\mathrm{e}^{-j\pi s^2/N}\mathrm{e}^{+j\pi(r-s)^2/N}.} \tag{30}$$

Define

$$u_k[r] = x[a+r]\mathrm{e}^{-j2\pi ar/N}\mathrm{e}^{-j\pi r^2/N}, \tag{31}$$

$$h[q] = \mathrm{e}^{+j\pi q^2/N}. \tag{32}$$

Substituting Eq. (30) into Eq. (6) yields

$$X_E[a+s] = \frac{N}{w}\mathrm{e}^{-j2\pi a^2/N}\mathrm{e}^{-j2\pi as/N}\mathrm{e}^{-j\pi s^2/N}\sum_{r=0}^{w-1} u_k[r]h[s-r]. \tag{33}$$

The summation is a linear convolution. Thus the original dense $w \times w$ EDPW block can be evaluated using chirp modulation, convolution, and output phase correction.

7.1. Equivalence theorem

Theorem. If the convolution in Eq. (33) is evaluated exactly, then the resulting Fast-EDPW coefficient is identical to the direct EDPW coefficient in Eq. (6) for every $x[n]$, N, w dividing N, block k, and $s \in \{0, \ldots, w-1\}$.

Proof. Equations (28) and (30) are algebraic identities. Substitution into the direct EDPW sum changes only the factorization and grouping of multiplicative terms; no term is omitted, approximated, or added. Equation (33) is therefore algebraically identical to Eq. (6). If

FFT/IFFT arithmetic is used to calculate the convolution, differences observed in finite-precision software arise from numerical roundoff rather than a change in the transform definition. □

Consequently,

$$\boxed{\boldsymbol{X}_{\text{FastEDPW}} = \boldsymbol{X}_{\text{EDPW}}} \tag{34}$$

in exact arithmetic, and

$$d(\boldsymbol{X}_F, \boldsymbol{X}_{\text{FastEDPW}}) = d(\boldsymbol{X}_F, \boldsymbol{X}_{\text{EDPW}}) \tag{35}$$

for any distance function d when exact arithmetic is assumed.

8. Fast EDPW Complexity and Crossover

For arbitrary integer w, the minimum radix-2 FFT length used for exact linear convolution is

$$\boxed{L(w) = 2^{\lceil \log_2(2w-1) \rceil}.} \tag{36}$$

Only when w is a power of two is this simplified to $L = 2w$. If the FFT of the fixed convolution kernel is precomputed, one block requires one forward L-point FFT, one inverse L-point FFT, and L pointwise complex products. Under the same radix-2 multiplication convention, the principal convolution-engine cost is

$$C_{\text{block,principal}} = L \log_2 L + L, \tag{37}$$

and for all N/w blocks,

$$\boxed{C_{\text{FastEDPW},\,principal} = \frac{N}{w}[L \log_2 L + L].} \tag{38}$$

For power-of-two w, Eq. (38) becomes

$$\boxed{C_{\text{FastEDPW},\,principal} = 2N[\log_2(2w) + 1].} \tag{39}$$

This establishes the $O(N \log w)$ principal-convolution order, but it does not yet represent the complete exact implementation cost.

8.1. Principal count versus optimized total estimate

Equation (38) excludes the block-dependent input and output phase/chirp products. With fixed N and w, invariant chirps and kernel transforms can be precomputed, and phase factors

can be combined. A useful optimized runtime estimate is therefore

$$\boxed{C_{\text{FastEDPW},\,opt} \approx \frac{N}{w}[L\log_2 L + L] + 2N.} \tag{40}$$

The added $2N$ is an implementation model, not a proven absolute lower bound. It represents approximately one input modulation and one output demodulation per coefficient after precomputation/fusion. Hardware-specific additions, memory traffic, scaling, and pipeline costs remain separate.

For $N = 2^p$ and power-of-two $w = 2^q$, the principal-only Fast-EDPW crossover against the radix-2 FFT is

$$2N(q+2) < \frac{N}{2}p \quad \Longleftrightarrow \quad \boxed{p > 4(q+2)}, \tag{41}$$

or equivalently $w < N^{1/4}/4$. Under the optimized-total estimate of Eq. (40), the stricter condition is

$$2N(q+3) < \frac{N}{2}p \quad \Longleftrightarrow \quad \boxed{p > 4(q+3)}, \tag{42}$$

or $w < N^{1/4}/8$. Thus the statement "full-spectrum Fast-EDPW is always more expensive than FFT" is not generally correct; it was true only for the previously examined N, w range.

8.2. Direct versus Fast realization

The word "Fast" denotes asymptotic acceleration of the dense EDPW block, not guaranteed superiority for every w. For power-of-two w, comparing $C_{\text{direct}} = Nw$ with Eq. (39) gives

$$w > 2[\log_2 w + 2] \tag{43}$$

as the principal-count region in which the chirp-convolution realization becomes cheaper. Under Eq. (40), the corresponding comparison is $w > 2[\log_2 w + 3]$. Consequently, very small windows are normally better evaluated directly, while larger windows favor Fast-EDPW.

9. Unified Complexity–Spectral Trade-off

For a fixed (N, w), the preferred EDPW realization is

$$\boxed{C_{\text{EDPW,best}}(N,w) = \min\{Nw,\ C_{\text{FastEDPW},opt}(N,w)\}.} \tag{44}$$

Table 1 places all three algorithms under one arithmetic convention. The Fast-EDPW column uses Eq. (40); the last column N/w is the approximate first-null width measured in global FFT bins and therefore warns against interpreting arithmetic economy as detailed spectral equivalence.

Table 1: Common complexity map: radix-2 FFT, Direct EDPW, and optimized-total Fast-EDPW estimate.

N	w	FFT	Direct	Fast opt.	Best/FFT	N/w
1,024	2	5,120	2,048	8,192	0.400	512
1,024	4	5,120	4,096	10,240	0.800	256
1,024	8	5,120	8,192	12,288	1.600	128
4,096	2	24,576	8,192	32,768	0.333	2,048
4,096	4	24,576	16,384	40,960	0.667	1,024
4,096	8	24,576	32,768	49,152	1.333	512
16,384	2	114,688	32,768	131,072	0.286	8,192
16,384	4	114,688	65,536	163,840	0.571	4,096
16,384	8	114,688	131,072	196,608	1.143	2,048

The table exposes the central design issue. Small w can produce a real arithmetic advantage, but the same choice broadens the local spectral response. Accordingly, three application-dependent window requirements should be distinguished:

$$\boxed{w_{\text{detect}} \leq w_{\text{depict}} \leq w_{\text{reconstruct}},} \tag{45}$$

where the terms denote, respectively, adequate threshold/occupied-region detection, useful coarse spectral morphology, and detailed FFT-like reconstruction. The inequalities express the usual information hierarchy rather than universal numerical thresholds; the required values depend on signal class, SNR, bandwidth, and decision rule.

Let $Q(N, w)$ be an application-specific quality metric and Q_{req} its acceptance threshold. The engineering choice is therefore

$$\boxed{w^* = \arg\min_w C_{\text{EDPW,best}}(N, w) \quad \text{subject to} \quad Q(N, w) \geq Q_{\text{req}}.} \tag{46}$$

Computational superiority is meaningful only relative to the spectral information required by the application. This joint cost–quality formulation is the appropriate interpretation of the Direct and Fast crossover results.

10. Selective Frequency-Region Processing

The comparison changes further when only a subset of EDPW blocks is required. Let B denote the number of selected blocks out of $K = N/w$. The principal Fast-EDPW cost is

$$C_{\text{FastEDPW},\, sel,principal} = B[L \log_2 L + L], \tag{47}$$

while an optimized-total estimate adds approximately $2Bw$ phase products. Selective Fast-EDPW is favorable against a complete FFT when the chosen accounting convention gives a

lower cost than $N \log_2 N/2$. For $N = 1024$, $w = 128$, $L = 256$, one selected block has 2,304 principal convolution products and approximately 2,560 products under the $+2w$ optimized estimate; two blocks require approximately 5,120 under that estimate, equal to the full FFT reference.

This motivates a workload-oriented architecture,

$$\boxed{\text{coarse scan} \;\rightarrow\; \text{threshold detection} \;\rightarrow\; \text{selective Fast-EDPW refinement},} \tag{48}$$

while making the accounting convention explicit. The practical value depends on occupancy, detection policy, required spectral quality, latency, and the cost of the coarse scanning stage.

11. Full-Spectrum and Selective Operating Regions

The revised analysis changes the earlier interpretation. Selective processing remains a major computational opportunity, but it is not the only one. Full-spectrum Direct EDPW can beat the radix-2 FFT for very small w according to Eq. (25), and full-spectrum Fast-EDPW can enter an arithmetic-win region at sufficiently large N according to Eq. (42). Neither result establishes spectral equivalence to the FFT. The correct comparison is therefore two-dimensional: arithmetic cost and task-required spectral information.

Consider a real-input analysis range from 0 to 2 GHz sampled at $F_s = 5$ GS/s. A 1 kHz global coordinate spacing corresponds to $N = 5,000,000$. The original EDPW partition requires w to divide N exactly, so the previous illustrative choice $w = 1024$ is replaced by

$$\boxed{w = 1000}, \qquad K = 5000. \tag{49}$$

One EDPW block then spans exactly 1 MHz on the global coordinate grid. The required linear-convolution FFT length is

$$L = 2^{\lceil \log_2(1999) \rceil} = 2048. \tag{50}$$

The principal selected-block cost remains

$$C_{\text{FastEDPW},block,principal} = 2048(11) + 2048 = 24,576, \tag{51}$$

and the optimized-total estimate is

$$C_{\text{FastEDPW},block,opt} \approx 24,576 + 2,000 = 26,576. \tag{52}$$

The radix-2-equivalent full-FFT reference is approximately 55.63×10^6 principal complex multiplications, so the corrected selected-workload arithmetic ratio is approximately 2.09×10^3.

This is an operation-count ratio, not a wall-clock speedup.

The local observation is now $T_w = 1000/(5 \times 10^9) = 200$ ns and the first-null scale is approximately $F_s/w = 5$ MHz. Thus the example deliberately demonstrates selective computation on a dense coordinate grid, not 1 kHz physical resolution.

If a lower-level architecture comparison is desired, time should be expressed in architecture cycles rather than seconds. Equal arithmetic lanes, multiplier/adder assumptions, memory ports, and pipeline assumptions are required for fairness.

12. Additional Theoretical Foundations

12.1. Frequency-grid theorem

Theorem 1 (global coordinate grid). For any EDPW coefficient whose kernel is evaluated as

$$X_E[m] = \alpha_w \sum_{n \in \mathcal{I}} x[n] \mathrm{e}^{-j2\pi nm/N_g}, \tag{53}$$

where N_g is the global kernel denominator and $\mathcal{I}$ contains w contributing samples, the coefficient indexed by m is evaluated at the Fourier coordinate

$$f_m = \frac{mF_s}{N_g}, \qquad \Delta f_{\text{grid}} = \frac{F_s}{N_g}. \tag{54}$$

The proof follows directly by equating the kernel $\mathrm{e}^{-j2\pi f_m n/F_s}$ with $\mathrm{e}^{-j2\pi mn/N_g}$. Hence N_g determines the frequency coordinates, whereas w determines how many observations contribute to each coefficient. These roles are mathematically distinct.

12.2. Single-tone response and finite-window selectivity

Let a complex sinusoid be $x[n] = A\mathrm{e}^{j(\omega_0 n+\phi)}$ and let an EDPW block begin at a. For $n = a+r$, $r = 0, \ldots, w-1$,

$$X_E[m] = \alpha_w A \mathrm{e}^{j\phi} \sum_{r=0}^{w-1} \mathrm{e}^{j\omega_0(a+r)} \mathrm{e}^{-j2\pi(a+r)m/N_g} \tag{55}$$

$$= \alpha_w A \mathrm{e}^{j\phi} \mathrm{e}^{ja\Delta\omega_m} \sum_{r=0}^{w-1} \mathrm{e}^{jr\Delta\omega_m}, \tag{56}$$

where $\Delta\omega_m = \omega_0 - 2\pi m/N_g$. The geometric sum gives

$$|X_E[m]| = |\alpha_w A| \left| \frac{\sin(w\Delta\omega_m/2)}{\sin(\Delta\omega_m/2)} \right|. \tag{57}$$

Thus the response is a w-sample Dirichlet kernel evaluated on an N_g-defined frequency grid. Its first zeros occur approximately at

$$|\Delta f| = \frac{F_s}{w}. \tag{58}$$

This proves the central distinction used throughout this paper:

$$\boxed{\Delta f_{\text{grid}} = F_s/N_g \quad \text{while} \quad \Delta f_{\text{null}} \approx F_s/w.} \tag{59}$$

A larger N_g samples the spectral response more densely; it does not by itself narrow the w-limited main lobe.

12.3. Detection is not spectral equality

For a requested frequency-index set $\mathcal{B}$ and threshold γ, define

$$D_{\mathcal{B}}(X) = \begin{cases} 1, & \max_{m\in\mathcal{B}} |X[m]|^2 > \gamma, \\ 0, & \text{otherwise.} \end{cases} \tag{60}$$

For a detection-oriented architecture, the task-level requirement can be

$$D_{\mathcal{B}}(X_E) = D_{\mathcal{B}}(X_F), \tag{61}$$

without requiring $X_E[m] = X_F[m]$ at every bin. Consequently, transform distance and detection agreement are separate performance measures. This distinction is fundamental to the original selective-search motivation of EDPW.

12.4. Selective-complexity crossover theorem

Let one selected Fast-EDPW block require C_b principal complex multiplications and let B blocks be evaluated. Then

$$C_{E,\text{sel}}(B) = BC_b. \tag{62}$$

Against a complete FFT cost C_F, selective Fast-EDPW is favorable under this arithmetic convention whenever

$$B < \frac{C_F}{C_b}, \tag{63}$$

so the largest integer number of strictly favorable blocks is

$$B_{\max} = \left\lceil \frac{C_F}{C_b} \right\rceil - 1. \tag{64}$$

If each selected block spans wF_s/N_g, Eq. (64) can be converted directly into a maximum favorable inspected bandwidth. The result is workload dependent and must include phase, memory, and control costs in a hardware-level study.

12.5. Exact fast/direct equivalence and error decomposition

The chirp-convolution factorization is an algebraic identity. In exact arithmetic,

$$\boldsymbol{X}_{\mathrm{FastEDPW}} = \boldsymbol{X}_E. \tag{65}$$

In finite precision write $\boldsymbol{X}_{\mathrm{FastEDPW}} = \boldsymbol{X}_E + \boldsymbol{\epsilon}_{\mathrm{num}}$. Therefore

$$\boldsymbol{X}_{\mathrm{FastEDPW}} - \boldsymbol{X}_F = \underbrace{(\boldsymbol{X}_E - \boldsymbol{X}_F)}_{\text{transform-definition difference}} + \underbrace{\boldsymbol{\epsilon}_{\mathrm{num}}}_{\text{implementation roundoff}}. \tag{66}$$

Numerical validation must measure these terms separately: direct EDPW versus Fast-EDPW tests the implementation, while EDPW versus FFT tests the transform definition.

12.6. Initial AWGN and coherent-processing implication

Let $x[n] = s[n] + v[n]$ with zero-mean white noise of variance σ_v^2. A coherent sum over w samples produces a deterministic sinusoidal contribution proportional to w, while the standard deviation of the summed white-noise component grows proportional to $\sqrt{w}$. Under consistent normalization, the coherent output SNR therefore scales proportionally to w. An N-sample DFT coefficient has the corresponding N-sample coherent-processing opportunity. Thus reducing the EDPW observation from N to w changes not only spectral selectivity but also noise-processing gain. Detection probability and false-alarm probability must therefore be evaluated as functions of both w and threshold γ.

12.7. Open theoretical question

The global-N_g phase structure permits a w-sample response to be evaluated on a much denser coordinate grid than F_s/w. Equation (57) proves that this does not by itself defeat the finite-window main-lobe scale. A remaining engineering question is whether the dense phase sampling nevertheless supplies useful sub-main-lobe *localization or detection* information for specified signal classes. This paper treats that question experimentally rather than assuming either outcome.

12.8. Computational frequency-grid limit

Let an abstract machine sustain at most X principal calculations in the allowed analysis interval. For a conventional FFT, the largest feasible transform length satisfies

$$\frac{N_{\max}}{2}\log_2 N_{\max} \le X, \tag{67}$$

and therefore the machine imposes the computational grid bound

$$\Delta f_{F,\min} = \frac{F_s}{N_{\max}}. \tag{68}$$

For selective Fast-EDPW, the local convolution engine is governed primarily by w and $L \simeq 2w$. Increasing a global grid parameter N_g changes chirp/phase coefficients and reduces F_s/N_g without requiring a local FFT of length N_g. However, if a *fixed physical bandwidth* must remain covered, the frequency span of one w-coefficient block, wF_s/N_g, shrinks as N_g increases; additional selected blocks may therefore be required. The relevant advantage is consequently not zero cost growth, but growth in small local-window increments rather than as a complete $N_g \log N_g$ transform.

13. Numerical Experiments

13.1. Experiment 0: introductory pure-sinusoid frequency localization

The introductory single-tone experiment uses

$$F_s = 1000 \text{ Hz}, \qquad N = 1024, \qquad f_c = 100 \text{ Hz}. \tag{69}$$

The global spacing is $\Delta f = 0.9765625$ Hz and the carrier is intentionally off-grid ($m_c = 102.4$). The conventional FFT peak occurs at 99.609375 Hz, giving a 0.390625 Hz bin error and normalized peak amplitude approximately 0.75634.

Table 2 summarizes the direct EDPW window sweep. The reported -3 dB widths are approximate discrete-spectrum measurements and are included as illustrative evidence rather than formal resolution claims. The trend is consistent with the theoretical distinction above: increasing w narrows the local spectral response while preserving the global index spacing.

Figure 1 displays the same result spectrally. The FFT is used as the full-record reference and the EDPW spectra are plotted for $w = 16, 32, 64, 128, 256$, and 512. The small-window cases distribute substantial spectral power over a wide neighborhood of the dominant frequency. As w increases, the EDPW response contracts toward the FFT reference. The figure is the visual counterpart of the F_s/w law: the arithmetic economy obtained by reducing w is accompanied

Table 2: Preliminary EDPW single-tone sweep, $F_s = 1000$ Hz, $N = 1024$, $f_c = 100$ Hz.

w	Peak (Hz)	Error (Hz)	Peak amplitude	Approx. -3 dB width (Hz)
16	101.562500	1.562500	1.022257	49.804688
32	99.609375	0.390625	0.985450	28.320312
64	99.609375	0.390625	1.013014	14.648438
128	99.609375	0.390625	1.000035	7.812500
256	99.609375	0.390625	0.985689	4.882812
512	99.609375	0.390625	0.937073	2.929688

by increasing spectral dispersion.

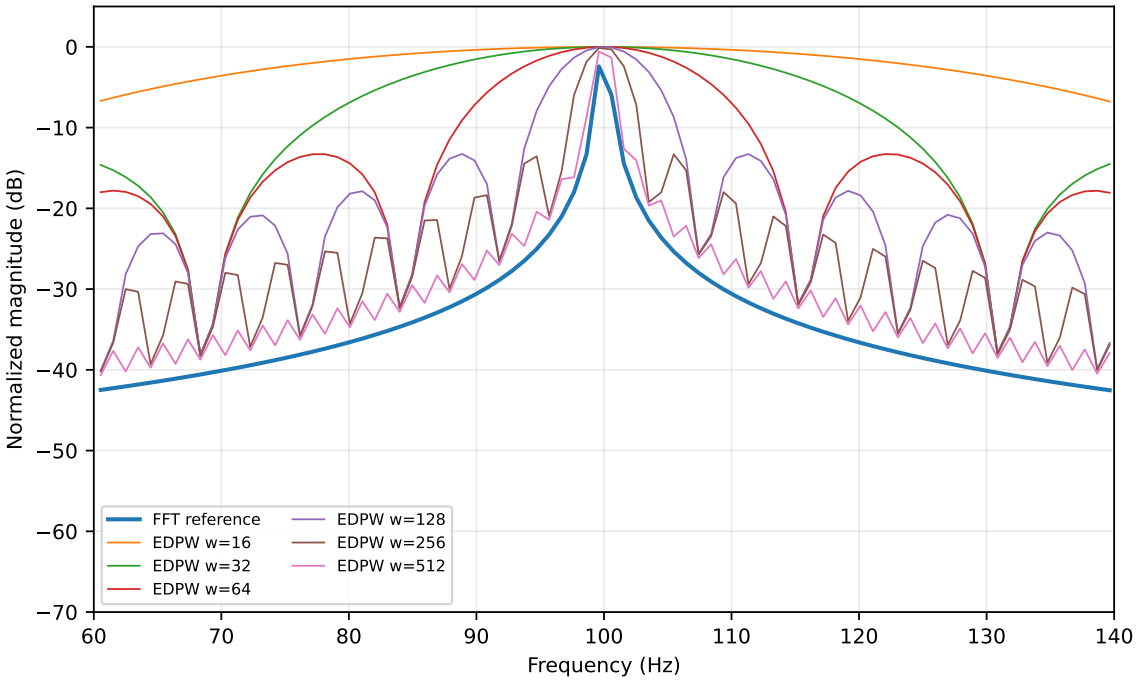


Figure 1: EDPW spectral dispersion versus window length for the 100 Hz single-tone experiment. Smaller w gives a broader response around the dominant frequency; increasing w progressively concentrates the spectrum toward the N-point FFT reference.

13.2. Window weighting as a dispersion-control trade-off

A natural question is whether filtering or window weighting can reduce the visible dispersion. Figure 2 tests this idea at $w = 64$ using rectangular, Hann, and Hamming weighting of the local EDPW samples with coherent-gain compensation. The weighted cases suppress parts of the distant sidelobe structure, but they also broaden the dominant main response. Consequently, conventional windowing can redistribute leakage and may be useful when sidelobe suppression is more important than narrow main-lobe localization, but it does not recover the spectral information lost by shortening the observation to w samples.

This distinction is fundamental. The finite observation establishes the basic information scale,

while a filter/window changes the trade-off within that scale. A future EDPW-specific filter may therefore be designed for a particular detector or interference environment, but it must be judged by task-level detection performance and computational cost rather than by visual smoothness alone.

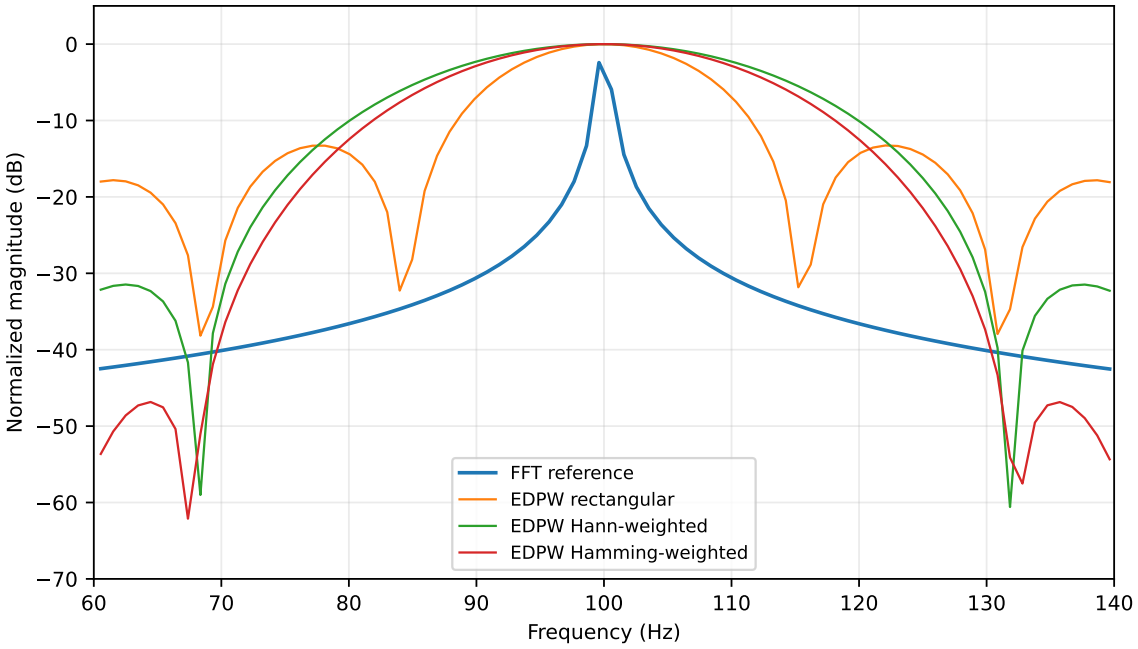


Figure 2: Exploratory local-window filtering at $w = 64$. Hann and Hamming weighting suppress portions of the sidelobe structure relative to the rectangular EDPW window, but broaden the main response. Filtering changes the leakage/selectivity trade-off; it does not recreate the full-record FFT resolution.

A separate implementation check is required for the fast formulation. The principal acceptance metric is

$$\epsilon_{\text{fast}} = \max_m |X_{\text{EDPW}}[m] - X_{\text{FastEDPW}}[m]|, \tag{70}$$

with relative forms added as needed. In double precision, errors near floating-point accumulation levels are expected if the indexing and convolution extraction are implemented correctly.

13.3. Experiment I: corrected selective wideband workload

The wideband experiment is retained, but the partition is corrected to $N = 5,000,000$ and $w = 1000$ so that $w \mid N$. Four deterministic tones occupy the requested 500–501 MHz region. The experiment asks whether the requested region can be evaluated selectively; it does not claim that the 1 kHz coefficient spacing provides 1 kHz physical resolution.

The earlier $w = 1024$ plots are not used as validation evidence in this analysis because 1024 does not exactly divide $5,000,000$ under the original EDPW definition. The companion MATLAB release regenerates the experiment with the admissible $w = 1000$ configuration.

Table 3: Experiment I corrected workload and arithmetic accounting.

Quantity	Value
Sampling rate F_s	5 GS/s
Acquired samples N	5,000,000
Global coordinate spacing F_s/N	1 kHz
EDPW window w	1,000
Block coordinate span	1 MHz
Local observation w/F_s	200 ns
First-null scale F_s/w	5 MHz
Convolution FFT length L	2,048
Fast-EDPW principal block cost	24,576
Fast-EDPW optimized block estimate	26,576
Full FFT reference	55.63×10^6
Corrected FFT/Fast ratio	$\approx 2.09 \times 10^3$

The scientific conclusion is unchanged: selective arithmetic can be extremely smaller than a complete transform for a narrow requested region, while the physical spectral selectivity remains governed by the short local observation.

Figure 3 makes that distinction visible. A deterministic complex band-limited test record was synthesized with nonzero spectral content only in the 500–501 MHz interval and processed by the full FFT, Direct EDPW, and Fast-EDPW. For visual comparison each curve is normalized to its own peak; this figure is intended to compare spectral *shape and localization*, not absolute gain. The FFT identifies the occupied interval sharply on the 1 kHz grid. With $w = 1000$, both EDPW realizations respond much more broadly because the local first-null scale is approximately 5 MHz. Direct EDPW and Fast-EDPW are visually coincident, as required by the equivalence theorem. Thus the selective calculation can detect activity in the requested region at very low arithmetic cost, but it does not preserve the FFT's sharp occupied-band edges at this w.

13.4. Experiment II: computational frequency-grid scalability at fixed machine capacity

Experiment II defines an abstract processing limit X rather than a processor clock. The reference machine is assumed capable of completing both FFT and Fast-EDPW arithmetic, but its maximum permitted work per analysis interval is set equal to the principal multiplication count of the $N = 5,000,000$ full FFT:

$$X = C_F(N) = \frac{N}{2}\log_2 N \approx 55.63 \times 10^6. \tag{71}$$

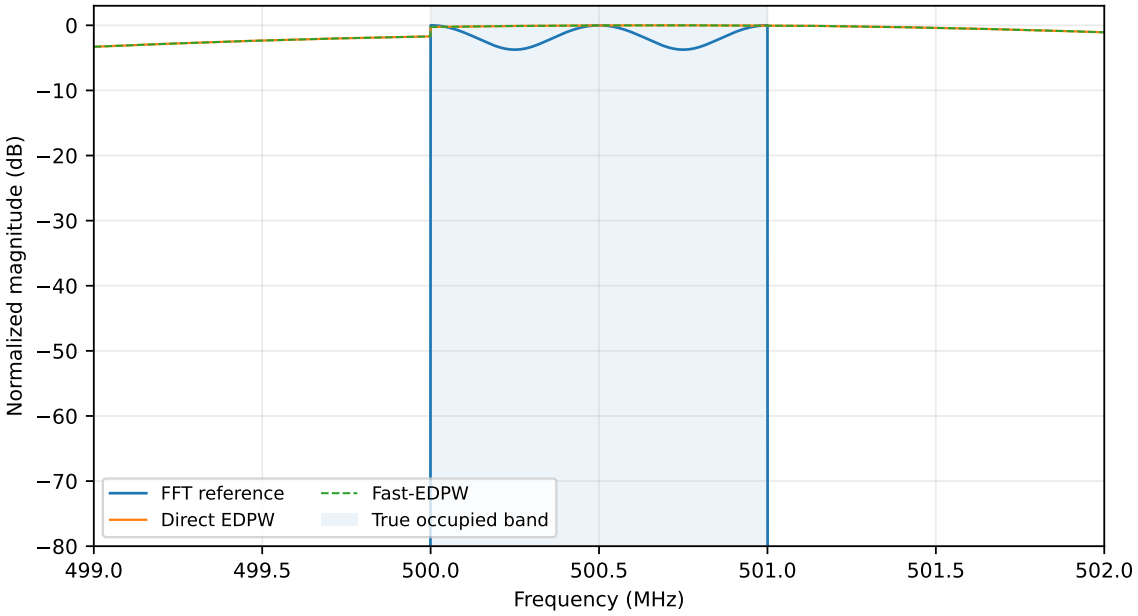


Figure 3: Experiment I, 500–501 MHz occupied-band example. The FFT reference sharply depicts the imposed band, whereas $w = 1000$ EDPW spreads the response over a much wider frequency region. Direct and Fast-EDPW overlap to numerical precision. Curves are individually peak-normalized to emphasize spectral shape.

At $F_s = 5$ GS/s this establishes the conventional FFT grid limit

$$\Delta f_F = \frac{F_s}{N} = 1 \text{ kHz}. \tag{72}$$

Now request a 50% larger global grid parameter,

$$N^* = 1.5N = 7,500,000, \tag{73}$$

which gives

$$\Delta f^* = \frac{5 \times 10^9}{7.5 \times 10^6} = 666.67 \text{ Hz}. \tag{74}$$

A complete conventional transform would require

$$C_F(N^*) = \frac{7.5 \times 10^6}{2} \log_2(7.5 \times 10^6) \approx 85.64 \times 10^6, \tag{75}$$

which is approximately 53.94% above the machine budget X and is therefore infeasible within the stipulated analysis interval.

For Fast-EDPW with $w = 1000$ and $L = 2048$, one selected block retains the principal cost

$$C_b = 2048[\log_2(2048) + 1] = 24,576. \tag{76}$$

The global denominator can be changed from N to N^* without enlarging the local $L = 2048$

FFT/IFFT engine. One qualification is essential: one $w = 1000$ block spans 1 MHz when $N = 5,000,000$, but only 666.67 kHz when $N^* = 7,500,000$. Covering the same complete 1 MHz physical region therefore requires two N^* blocks in the simple non-overlapping selection used here, for a principal cost of only

$$C_{E,1\mathrm{MHz}}(N^*) = 2C_b = 49,152. \tag{77}$$

This remains orders of magnitude below the unavailable 85.64-million-product full FFT. The experiment therefore tests a computational decoupling: finer global grid coordinates can be requested without replacing the local Fast-EDPW engine by an N^*-point FFT.

A deliberately off-grid sinusoid at $f_0 = 500.5004$ MHz can be evaluated on the 1 kHz N grid and on the 666.67 Hz N^* coordinate grid while retaining $w = 1000$. The denser coordinate sampling is interpreted only as grid refinement, not as a narrower physical main lobe.

Figure 4 shows the corresponding $N = 5,000,000$ spectral comparison at the original 1 kHz grid. The FFT and both EDPW realizations identify the nearest coefficient at 500.500 MHz, giving a 0.4 kHz grid-localization error for this test. The important difference is the spectral shape: the FFT concentrates the tone into its narrow leakage pattern, while $w = 1000$ EDPW produces the expected broad local response. Direct and Fast-EDPW again coincide. This figure therefore separates *frequency-grid localization* from *physical spectral selectivity*: a coefficient can lie close to f_0 even though the finite-w response is several megahertz wide.

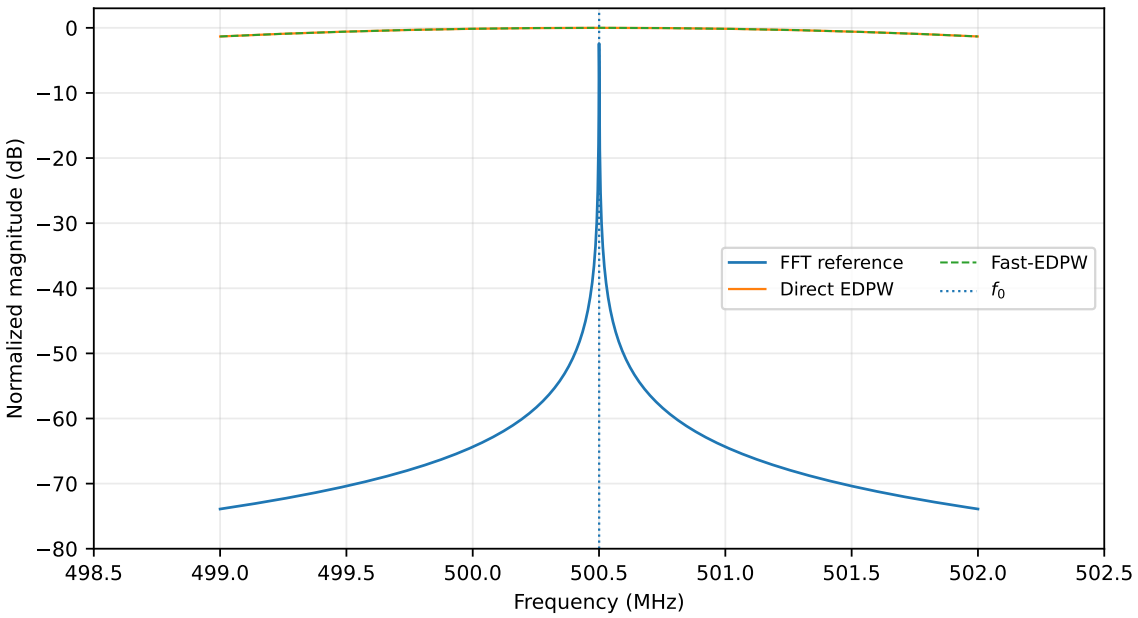


Figure 4: Experiment II, off-grid sinusoid at $f_0 = 500.5004$ MHz. FFT, Direct EDPW, and Fast-EDPW use the same acquired signal. The EDPW curves overlap, while their broad response illustrates the F_s/w selectivity limit despite fine global frequency labels.

The experiment does *not* claim 666.67 Hz two-tone resolving power for $w = 1000$. Equation (57) predicts a first-null scale near $F_s/w = 5$ MHz. The result being tested is computa-

Table 4: Experiment II: fixed-X computational frequency-grid comparison.

Method	Grid parameter	Grid spacing	Principal products
FFT at machine limit	5,000,000	1,000 Hz	55.63 M
FFT requested at $1.5N$	7,500,000	666.67 Hz	85.64 M (infeasible)
Fast-EDPW, one block	5,000,000	1,000 Hz	24,576
Fast-EDPW, same 1 MHz coverage at $1.5N$	7,500,000	666.67 Hz	49,152

tional: the conventional full FFT cannot refine its grid beyond the fixed-X boundary without increasing transform work, whereas selective Fast-EDPW can refine its global coordinate grid with comparatively small local-window cost. Whether those extra grid samples improve detection or localization is a separate signal-information question and must be assessed separately by task-level metrics.

The N^* construction is retained only as a computational-coordinate experiment. It does not create additional observations, and therefore it does not establish improved physical two-tone resolution. Any generalized N^* implementation must also define coefficient indexing and block pairing consistently over the enlarged coordinate grid; the present result should not be read as a completed generalized transform definition.

14. Comparative Spectral Algorithms: Mathematical Background

The methods compared here do not all compute the same output. An exact full-record DFT, an exact selected-bin DFT, a local-window estimator, and a time–frequency representation are distinct mathematical products. Comparisons are consequently separated into *full-spectrum depiction*, *selected-bin or selected-band evaluation*, and *time-localized analysis*. All methods use the same acquired record when the task requires the same input information.

14.1. Full DFT and radix-2 FFT

For N observations, the unnormalized DFT is $X[k] = \sum_{n=0}^{N-1} x[n]e^{-j2\pi kn/N}$. The radix-2 FFT reorganizes this exact sum using butterfly stages, with the conventional principal reference of $\frac{N}{2}\log_2 N$ complex multiplications and $N\log_2 N$ complex additions for N a power of two [3]. These are conventional reference counts, not optimized hardware instruction counts. The FFT reproduces the finite-record DFT; it cannot defeat finite-observation leakage or noise.

14.2. Direct EDPW and exact Fast-EDPW

For $I_b = \{bw, \ldots, (b+1)w - 1\}$ and $k \in I_b$, the EDPW output is

$$X_E[k] = \frac{N}{w} \sum_{n \in I_b} x[n] e^{-j2\pi kn/N}, \qquad N/w \in \mathbb{N}. \tag{78}$$

This is not an w-point FFT: the phase denominator is N. Direct evaluation uses Nw kernel products. Fast-EDPW implements the *same coefficients* through an exact chirp-convolution factorization, related algebraically to the chirp-convolution techniques of Bluestein [4]. Let $L = 2^{\lceil \log_2(2w-1) \rceil}$. With the kernel FFT reusable, its principal convolution multiplication count is $(N/w)(L \log_2 L + L)$. Input/output chirp and phase multiplications must be counted separately. The approximate optimized phase-inclusive count is the principal count plus $2N$; this is an implementation model rather than an exact universal total.

14.3. Output-pruned FFT

A pruned FFT evaluates only a requested set $\mathcal{K}$ of the *same* N-point DFT coefficients, omitting unnecessary branches of a butterfly network [6, 8]. Its cost depends on N, $|\mathcal{K}|$, the locations of selected indices, the radix and pruning strategy. Thus no single $O(|\mathcal{K}| \log N)$ expression is valid as an exact universal operation count. We report actual butterfly counts from the specified output-pruned radix-2 implementation and compare its outputs against the full DFT.

14.4. Goertzel recurrence

For a selected DFT bin k with $\omega_k = 2\pi k/N$, define $s_{-1} = s_{-2} = 0$ and

$$s_n = x[n] + 2\cos(\omega_k) s_{n-1} - s_{n-2}, \quad n = 0, \ldots, N-1. \tag{79}$$

The selected coefficient follows from $X[k] = e^{-j\omega_k(N-1)}(s_{N-1} - e^{-j\omega_k} s_{N-2})$ (with indexing interpreted consistently). Goertzel evaluates M selected coefficients using $O(NM)$ scalar recurrences, with constants depending on whether complex input, real coefficients and final complex operations are separately counted [1]. It is an exact selected-DFT method in exact arithmetic.

14.5. Chirp-Z transform (CZT)

The CZT computes $Y[k] = \sum_{n=0}^{N-1} x[n] A^{-n} W^{nk}$, $k = 0, \ldots, M-1$. The DFT is recovered by $A = 1, W = e^{-j2\pi/N}, M = N$. Bluestein-type convolution uses an FFT length $L_C \geq N+M-1$, typically a supported power of two, and requires $O(L_C \log L_C)$ operations plus chirp generation

and multiplication [5, 4]. CZT is valuable for chosen frequency grids, but a zoomed grid does not itself improve the information-limited resolution of a fixed record.

14.6. Short-time Fourier transform (STFT)

For analysis window $g[r]$, length w, hop H and frame index b, the STFT is $S[b,k] = \sum_{r=0}^{w-1} x[bH+r]g[r]e^{-j2\pi kr/w}$. It generates a time–frequency representation, not the same output as the full-record DFT [2, 7]. If B frames are analyzed using w-point FFTs, the conventional principal multiplication reference is $B(w/2)\log_2 w$, plus windowing and overlap/data movement. STFT comparisons therefore use time-localization tasks rather than ranking its bins against the N-point FFT as though they were equivalent.

15. Expanded Total Computational Complexity Comparison

Table 5 records a common set of analytical baselines. Multiplication types and setup costs must not be conflated; in particular Goertzel recurrences with real coefficients are not equivalent to full complex multiplications. FFT-based convolution formulas exclude implementation-specific FFT planning and memory traffic. Here $M = |\mathcal{K}|$, B is the number of STFT frames, $L = 2^{\lceil \log_2(2w-1) \rceil}$, and $L_C = 2^{\lceil \log_2(N+M-1) \rceil}$.

Table 5: Expanded computational comparison. Expressions are analytical references; counts use explicitly different arithmetic units where indicated.

Method	Multiplication / arithmetic reference	Output and caveat
Radix-2 FFT	$(N/2)\log_2 N$ complex mult.; $N\log_2 N$ complex adds	Exact complete N-point DFT
Direct EDPW	Nw complex kernel products; $N(w-1)$ adds	Local-support approximation; full EDPW grid
Fast-EDPW	$(N/w)(L\log_2 L + L)$ principal complex mult.; $+2N$ estimated optimized phase products	Exact EDPW coefficients; not exact DFT generally
Output-pruned FFT	Actual butterfly counts from selected-bin radix-2 recursion	Exact selected DFT; selection dependent
Direct selected DFT	NM complex products; $(N-1)M$ adds	Exact selected DFT; validation oracle
Goertzel	NM second-order real-coefficient recurrences + final complex steps	Exact selected DFT; count separately by arithmetic type
CZT	$O(L_C \log_2 L_C)$ convolution work + chirps	Arbitrary frequency-grid evaluation
STFT	$B(w/2)\log_2 w$ principal complex FFT mult. + window products	Time–frequency output; distinct task

16. Experiment III: Comparative Spectral Vulnerability

The campaign compares exact full-record spectral methods against EDPW's local-observation coefficients. A common $N = 1024$ complex-sample record at $F_s = 4096$ Hz contains tones at 311.25, 316.5, and 927.8 Hz with amplitudes 1, 0.30, and 0.15. For the numerical table and figures below, the complex perturbation is explicitly deterministic:

$$\eta[n] = \frac{0.025}{\sqrt{2}} \left[\sin(2\pi\sqrt{2}\, n) + j \cos(2\pi\sqrt{3}\, n) \right]. \tag{80}$$

This deterministic perturbation is *not* a sample of complex AWGN; separate stochastic scripts are reserved for Monte Carlo studies. The results in this section were independently calculated in Python/NumPy from the stated equations. The companion MATLAB publication-results script reproduces the same input and output definitions, but MATLAB R2016a execution remains to be confirmed.

For each w, define $D_{\text{mag}} = \| \, |X_E| - |X_F| \, \|_2 / \| \, |X_F| \, \|_2$ and Pearson correlation ρ of the two magnitude vectors. The peak-location error refers to the global maximum over nonnegative frequencies relative to the strong tone at 311.25 Hz; it is *not* a measure of weak-tone detectability. The cold-call Fast-EDPW column is an estimated complex-multiplication reference including one kernel FFT and $2N$ optimized phase products; it is not measured runtime.

Table 6: Computed spectral comparison on the common deterministic record. Direct EDPW and Fast-EDPW have identical ideal outputs; operation counts are model-based.

w	ρ	D_{mag}	Peak err. (Hz)	Direct products	Fast cold est.
8	0.2227	12.5683	3.25	8,192	12,320
16	0.2955	8.8909	0.75	16,384	14,416
32	0.3850	6.2136	0.75	32,768	16,576
64	0.4971	4.1601	0.75	65,536	18,880
128	0.6327	2.7752	0.75	131,072	21,504
256	0.7851	1.5031	0.75	262,144	24,832
1024	1.0000	0.0000	0.75	1,048,576	37,888

16.1. Common-bin comparison across the spectral methods

A separate, reproducible comparison evaluates exactly the same $N = 1024$ complex record as above at the predetermined $M = 33$ DFT bins $k = 64, \ldots, 96$ (frequency range 256–384 Hz). The interval is selected *before* inspecting the test signal; it is not a blind full-spectrum discovery task. Full FFT, output-pruned FFT, Goertzel and a DFT-grid CZT all compute the same full-record coefficients, to numerical precision. Direct and Fast-EDPW compute identical *EDPW* coefficients at $w = 64$, which are not generally equal to the DFT. All

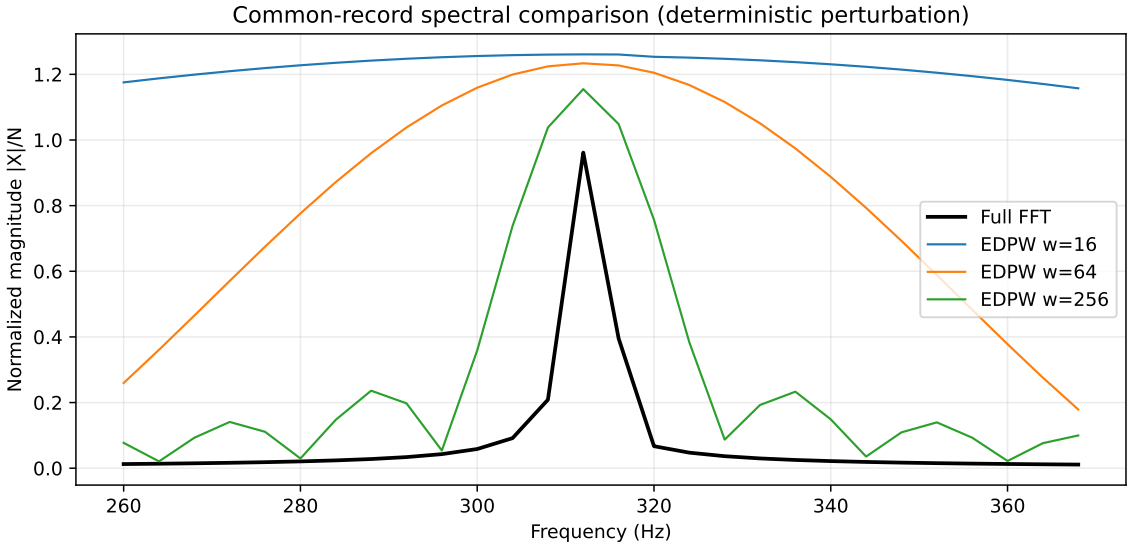


Figure 5: FFT reference and selected EDPW spectra from exactly the same input. A short local observation disperses energy across frequency.

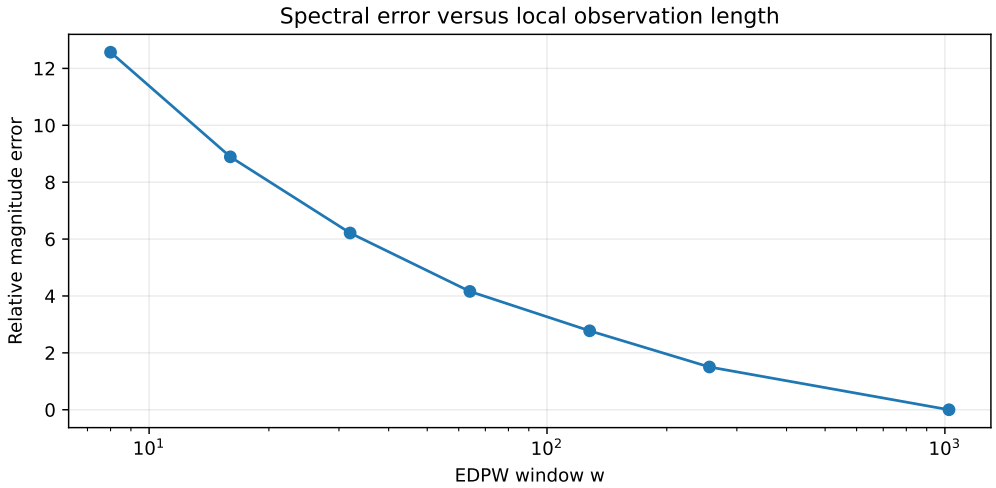


Figure 6: Full-spectrum relative magnitude error versus EDPW window size. The $w = N$ endpoint is the DFT limiting case.

methods have access to the same acquired input record; EDPW intentionally uses only the paired local sample groups for its individual outputs.

Table 7: Common-bin experiment, independently calculated from the stated equations in NumPy. Errors are relative complex ℓ_2 distances to the 33 FFT coefficients; peak errors refer to the strongest 311.25 Hz component. Counts are arithmetic references, *not* measured timings. Goertzel counts real-coefficient recurrences, not complex products.

Method	Rel. error	Peak err. (Hz)	Count	Arithmetic unit
Full FFT	0	0.75	5,120	complex products
Output-pruned FFT	$6.78 \cdot 10^{-16}$	0.75	6,143	complex products
Goertzel	$5.81 \cdot 10^{-14}$	0.75	33,792	real recurrences
CZT (DFT grid)	$2.84 \cdot 10^{-13}$	0.75	36,897	complex products
Direct EDPW ($w = 64$)	4.1730	0.75	65,536	complex products
Fast-EDPW ($w = 64$)	4.1730	0.75	18,880	complex products

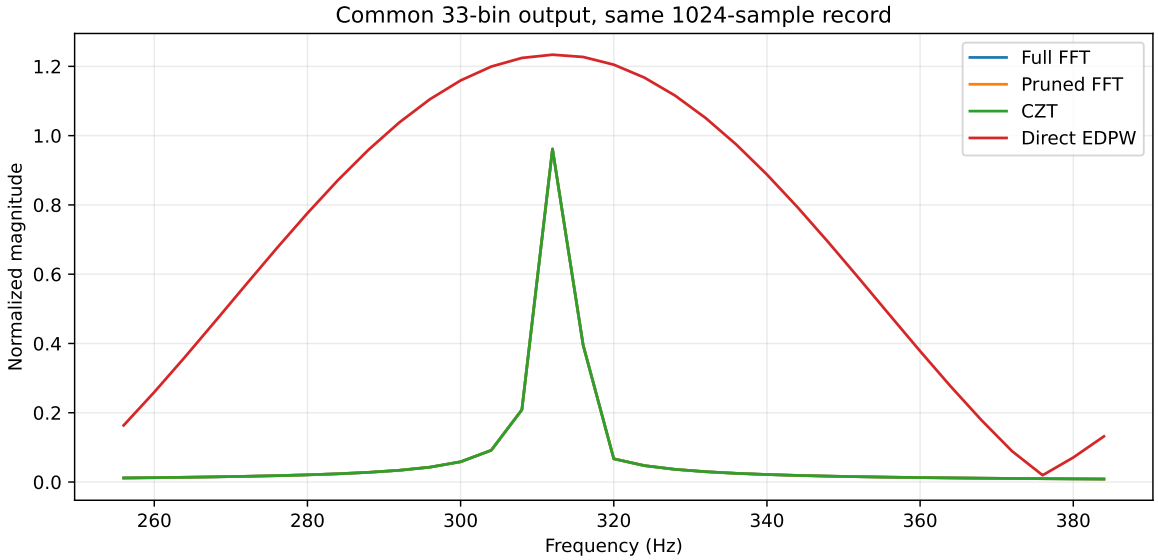


Figure 7: Comparison on the same 33 selected frequency bins. Exact DFT methods overlap almost completely; the plotted EDPW curve differs because it uses local observation windows. The Fast-EDPW spectrum coincides with Direct EDPW and is not separately overplotted.

The operation counts correspond to the *implemented workloads*: the FFT computes all N bins, while pruned FFT, Goertzel and CZT compute 33 requested bins. The two EDPW counts above represent a *full* EDPW transform because the accompanying general-purpose helpers compute the complete output; they must not be interpreted as selective-block counts. Pruned FFT is more expensive than the complete radix-2 FFT under the specified unoptimized butterfly-count convention for this relatively broad selected interval. CZT includes a cold-call chirp-kernel transform and chirp multiplications. Different arithmetic types, FFT planning, memory movement, and actual runtime are not made equivalent by the count column. The table therefore demonstrates accuracy and accounting, *not* a universal speed ranking.

16.2. STFT: a separate time-localization experiment

The STFT does not return the same product as the preceding methods. To assess its intended use, a second input contains a 320 Hz complex sinusoid active only for sample indices $384 \leq n < 640$ of a 1024-sample record. A Hann-window STFT with window length 128 and hop 32 produces the time–frequency representation in Fig. 8. The gated tone is localized in time, with the expected temporal smearing from the analysis window. This figure is *not* a claim that STFT reproduces the 1024-point DFT bins or that it outperforms the FFT in general.

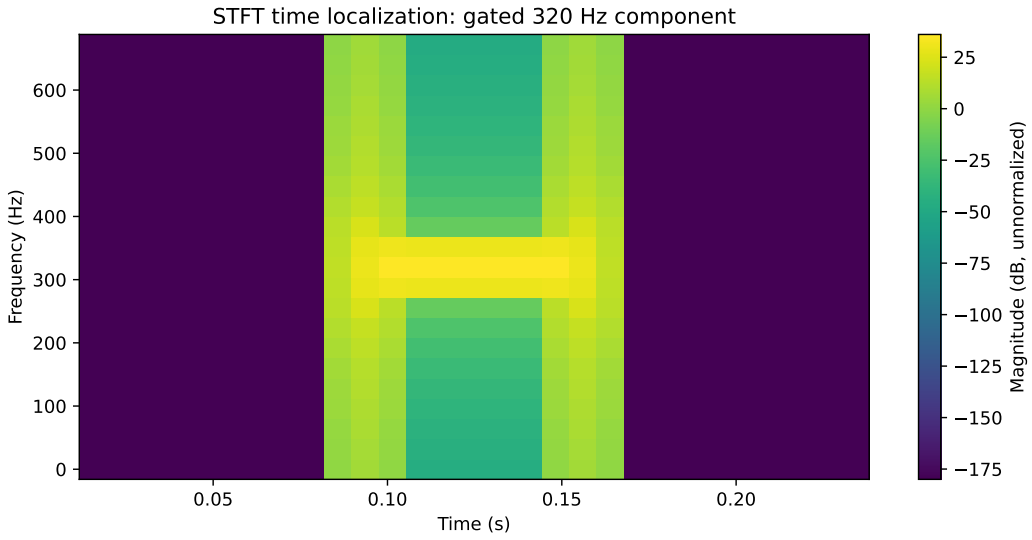


Figure 8: STFT on a separate gated-tone localization task, $F_s = 4096$ Hz, Hann window 128, hop 32. The active interval is 0.09375–0.15625 s; window overlap creates expected boundary smearing.

Verification scope. The figures and table were generated using the included independent NumPy reference program. The public MATLAB experiment implements the same specified scenarios and writes a comparison CSV and figures; its execution in MATLAB R2016a has not yet been verified. This is a deterministic accuracy demonstration, not a measured throughput or Monte Carlo detection-probability study.

The table and figures illustrate the cost–fidelity trade-off, not a universal ranking. Exact selected-bin methods (pruned FFT, Goertzel and DFT-grid CZT) target the full-record DFT and therefore share its coefficients, up to floating-point error, at matching frequency bins. They are not interchangeable with local EDPW output when $w < N$. In particular, a large magnitude-spectrum error at short w is compatible with a low arithmetic count. The close weaker tone cannot be declared detected solely because the strongest peak is localized. A fair detection comparison must also charge discovery of unknown regions and evaluate misses and false alarms at matched operating points. The separate Monte Carlo and fixed-band

drivers are provided for that investigation, without asserting unverified probability estimates.

16.3. Additional evidence: spectral fidelity, arithmetic cost, and blind discovery

The following three figures extend the comparative experiment beyond the 33-bin spectral overlay. They were computed independently in NumPy using the same stated mathematical definitions, with deterministic inputs so that the plotted values can be reproduced without dependence on a particular random-number generator. The companion MATLAB driver uses the same inputs and metrics. These are *model-based arithmetic comparisons*, not wall-clock measurements or demonstrations of MATLAB execution speed.

Spectral fidelity versus computational cost. Figure 9 uses the same $N = 1024$, $F_s = 4096$ Hz deterministic composite record as the preceding Section 16 comparison. For each divisor w of N, it displays the magnitude-spectrum error and correlation against the full FFT, together with the Direct and Fast-EDPW estimated complex-product counts normalized by the radix-2 FFT reference. The full-spectrum FFT cost is $N \log_2(N)/2 = 5120$ complex products under the stated counting convention. The short-window EDPW spectra exhibit large magnitude error even when their arithmetic estimates are relatively low. As w approaches N, fidelity improves, but the local-observation economy is progressively lost. Warm-kernel estimates assume a reusable precomputed convolution kernel; cold-call estimates include its FFT. The phase-product term is an optimized estimate and excludes coefficient generation and memory traffic. The figure therefore illustrates a *conditional cost–fidelity trade-off*, not a runtime speedup or an equal-accuracy advantage.

Quality-conditioned choice. Figure 10 minimizes the estimated full-spectrum cost over Direct EDPW, cold-call Fast-EDPW, and the FFT, subject to an upper bound on relative magnitude error. The FFT is always eligible because its error relative to itself is zero. At stringent tolerances it is therefore the least expensive eligible option for this record under the stated cost model; only after allowing substantial spectral mismatch can a cheaper local-window configuration become eligible. The step-like curve reflects the discrete tested window sizes, not a continuous optimization law. In particular, these thresholds are *specific to the illustrated signal* and do not guarantee detection, weak-tone recovery, or fidelity for other inputs.

Blind coarse-band discovery. Figure 11 addresses the region-selection issue using a *different* deterministic $N = 256$ record, sampled at $F_s = 1024$ Hz, with tones at 117.3, 309.6, and 443.2 Hz (amplitudes 1, 0.45, and 0.22). The spectrum is partitioned into 16 predetermined coarse bands. Each method scores every band by the sum of squared magnitudes of its own full-spectrum coefficients; no method receives the tone locations

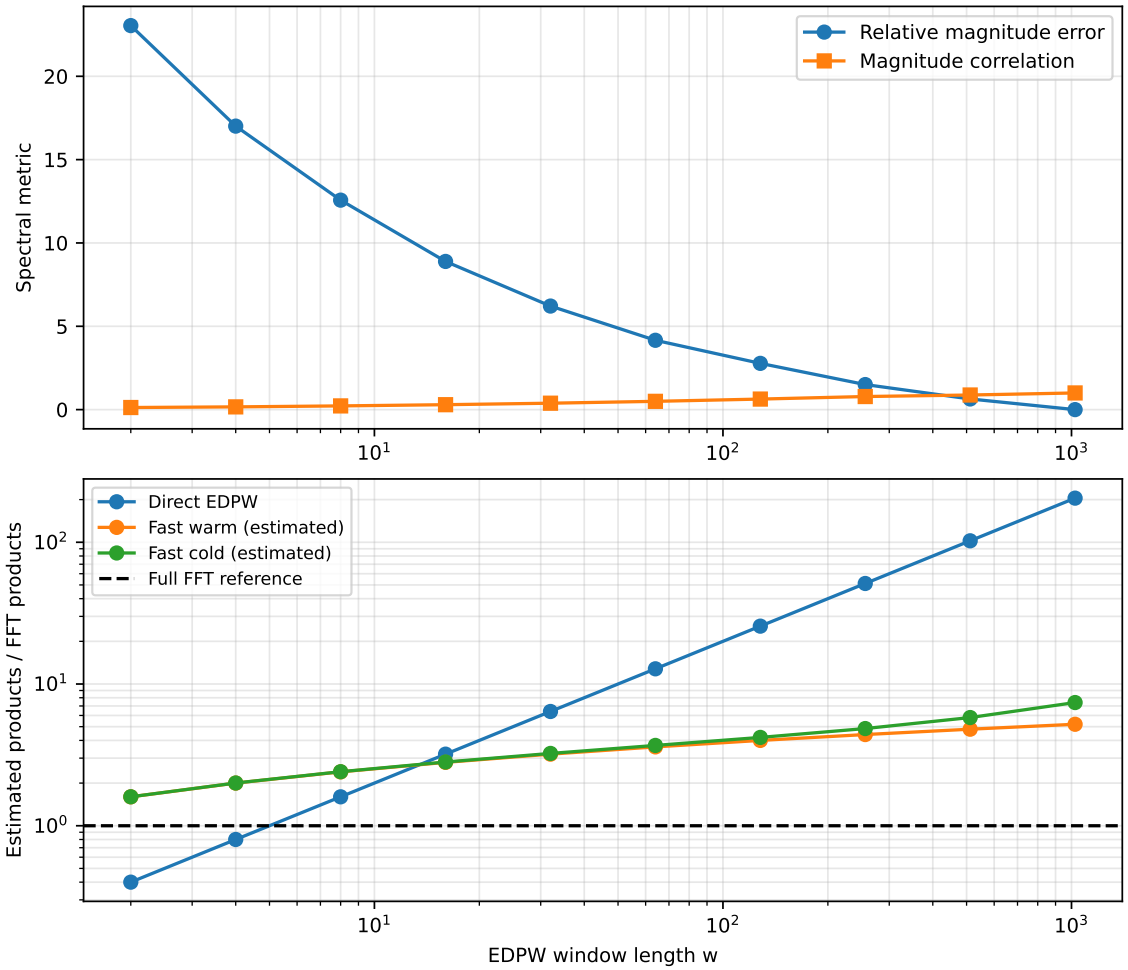


Figure 9: Magnitude-spectrum fidelity and estimated complex-multiplication cost for Direct and Fast-EDPW on the common deterministic $N = 1024$ record. Counts are normalized to the full radix-2 FFT reference; warm and cold Fast-EDPW estimates differ by the kernel FFT preparation cost.

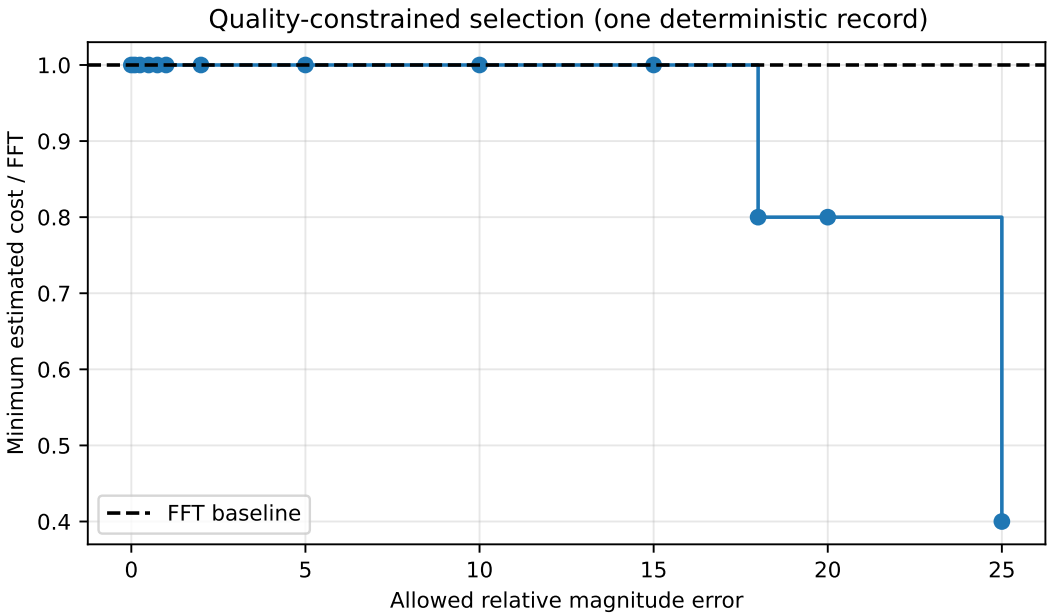


Figure 10: Minimum estimated arithmetic cost subject to a permitted relative magnitude-spectrum error on the deterministic record. The full FFT is an explicit eligible candidate at every threshold; costs are not measured runtimes.

beforehand. The shaded bands indicate the true occupied regions *for scoring only*. Since EDPW modifies the distribution of spectral energy, the relative band scores and their rankings may differ from the FFT even though the same input was acquired. This single-record illustration does not establish detection probabilities or false-alarm rates; those require a properly controlled Monte Carlo campaign with a defined detection rule.

The three additions reinforce distinct conclusions: cost alone does not imply equivalent spectral information; accuracy requirements must be included when selecting an implementation; and a frequency-selective processing policy must account for discovery of unknown occupied regions. Their reproducible deterministic construction is documented in the companion code, while MATLAB R2016a execution remains an independent validation step.

17. Proposed Distance Metrics for the Numerical Campaign

For every tested w, the following metrics will be calculated.

17.1. Complex spectral distance

$$D_2(w) = \frac{\|\boldsymbol{X}_{\text{FastEDPW}} - \boldsymbol{X}_F\|_2}{\|\boldsymbol{X}_F\|_2}. \tag{81}$$

This metric is phase sensitive.

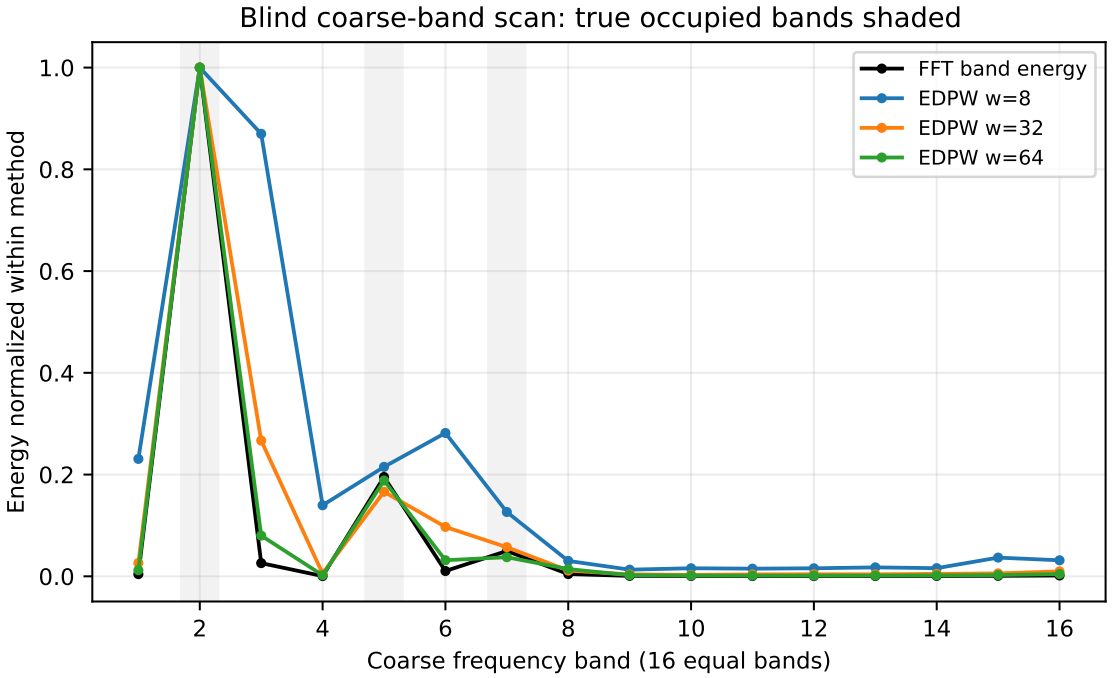


Figure 11: Blind 16-band energy scan on a separate deterministic $N = 256$ record. Each curve is normalized by its own maximum band score for shape comparison; gray shading identifies the true tone-containing bands, which are not supplied to the algorithms during scanning.

17.2. Magnitude-spectrum distance

$$D_{\mathrm{mag}}(w) = \frac{\| \, |\boldsymbol{X}_{\mathrm{FastEDPW}}| - |\boldsymbol{X}_F| \, \|_2}{\| \, |\boldsymbol{X}_F| \, \|_2}. \tag{82}$$

This is often more directly relevant to spectral depiction and threshold detection.

17.3. Maximum normalized bin error

$$D_\infty(w) = \frac{\max_m |X_{\mathrm{FastEDPW}}[m] - X_F[m]|}{\max_m |X_F[m]|}. \tag{83}$$

17.4. Task-level metrics

Transform distances alone do not determine engineering usefulness. The validation should also record peak-frequency error, peak-amplitude error, two-tone separability, threshold-crossing agreement, false-positive/false-negative detection rates under noise, and computational work as a function of the fraction of spectrum inspected.

18. Experimental Status and Remaining Engineering Work

The public MATLAB package contains R2016a-targeted drivers for Experiments I and II and the Experiment III spectral-vulnerability, common-bin, time-localization, arithmetic-budget, Monte Carlo, blind-scan, fixed-band and quality-conditioned drivers. Their existence establishes *implementation coverage*, not successful execution or verified scientific outcomes. The runner executes an initial self-test, records per-experiment status, and is designed to save figures and workspace data. The scripts must be executed in MATLAB R2016a, and their output must be inspected before numerical conclusions are accepted. A true matched-budget comparison requires matching useful task outputs, not merely reporting whether an algorithm falls below an operation threshold. Memory accounting, blind region-discovery cost, and detection operating characteristics remain open. Existing Experiment I and II manuscript material is retained as prior baseline evidence, with its stated scope and limitations.

The principal interpretive rule is

$$\boxed{\text{arithmetic economy} \neq \text{spectral fidelity} \neq \text{detection sensitivity.}} \tag{84}$$

No algorithm is declared superior from asymptotic order alone. The final comparison will report the actual computed output, input information used, arithmetic counting assumptions, and observed numerical fidelity.

19. Limitations and Claims Deliberately Not Made

This study does not claim that EDPW is a universally faster FFT, that the global spacing F_s/N equals EDPW physical resolution for $w < N$, or that the present chirp-convolution implementation is arithmetically optimal. The comparative Experiment III is not yet a completed Monte Carlo campaign, and no generalized detection-probability claims are made. No hardware throughput claim is made.

The current theoretical claims are narrower and stronger: the EDPW transform is explicitly defined; its exact coefficient difference from the DFT is known; its normalized Frobenius operator distance from the DFT is known; the DFT is recovered exactly at $w = N$; and the chirp-convolution Fast-EDPW is algebraically equivalent to direct EDPW.

20. Conclusions

The revised analysis replaces a simple "FFT for full spectrum, EDPW for selective spectrum" interpretation with a more precise operating map. Direct full-spectrum EDPW has a genuine arithmetic-win region, $w < \frac{1}{2}\log_2 N$, although that region is spectrally coarse. Fast-EDPW

reduces the dense EDPW block from $O(Nw)$ to $O(N \log w)$ principal convolution work and can itself enter a full-spectrum arithmetic-win region for sufficiently large N and small w. Complete implementation accounting is stricter because chirp/phase products, additions, memory traffic, and hardware organization also matter.

The central design variable is therefore not merely the algorithm name but the pair (N, w) together with the required spectral task. The preferred EDPW implementation is the cheaper of Direct and Fast realization, and w should be chosen by minimizing that cost subject to a detection, depiction, or reconstruction-quality requirement. The expanded complexity table makes the computational crossovers visible; the N/w selectivity column prevents those crossovers from being mistaken for spectral equivalence.

Selective processing remains especially important because independent EDPW blocks permit computation to be omitted outside requested regions. The corrected wideband example uses $N = 5,000,000$ and the admissible $w = 1000$, preserving the large selected-workload arithmetic reduction while making explicit that the 1 kHz coordinate spacing coexists with an approximately 5 MHz local first-null scale.

Experiments I–III and the crossover/fidelity campaign separate representation quality, computational architecture, and task-level detection performance. The comparison distinguishes complete-DFT algorithms from local-observation estimators: arithmetic savings cannot be interpreted as equal spectral fidelity without matching the requested information task. Accordingly, the defensible claim of this paper is conditional rather than universal: EDPW can provide substantial arithmetic savings when the required information can be obtained from its local-window coefficients, and the appropriate w and realization must be selected jointly from cost and quality requirements.

The same mechanism that creates the computational economy also defines its price. Reducing w reduces the number of processed products, but shortens the effective observation and therefore disperses spectral power over a wider neighborhood of a dominant component. This dispersion is not an implementation defect of Fast-EDPW; it is already present in Direct EDPW and follows from the finite-window response. Fast spectral identification is therefore obtained by accepting a coarser spectral representation when the task permits it.

This trade-off also suggests a natural two-stage use of EDPW. A small or moderate w can act as a preliminary/coarse spectral-analysis tool that scans economically for occupied or significant frequency regions. Once a region of interest is detected, the analysis can zoom into that region with a larger w, a focused Fast-EDPW evaluation, or another higher-resolution method such as an FFT/CZT stage. In this interpretation EDPW need not replace the FFT as a complete spectral reconstructor; it can reduce the amount of spectrum that requires expensive detailed analysis. The resulting architecture is therefore *scan* $\rightarrow$ *detect* $\rightarrow$ *focus/zoom*, with computational economy and spectral dispersion treated as an explicit engineering exchange rather than independent claims.

21. References

References

[1] G. Goertzel, “An Algorithm for the Evaluation of Finite Trigonometric Series,” *American Mathematical Monthly*, vol. 65, no. 1, pp. 34–35, 1958.

[2] J. B. Allen, “Short Term Spectral Analysis, Synthesis, and Modification by Discrete Fourier Transform,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 3, pp. 235–238, 1977.

[3] J. W. Cooley and J. W. Tukey, “An Algorithm for the Machine Calculation of Complex Fourier Series,” *Mathematics of Computation*, vol. 19, no. 90, pp. 297–301, 1965. doi:10.1090/S0025-5718-1965-0178586-1.

[4] L. I. Bluestein, “A Linear Filtering Approach to the Computation of the Discrete Fourier Transform,” *IEEE Transactions on Audio and Electroacoustics*, vol. 18, no. 4, pp. 451–455, 1970. doi:10.1109/TAU.1970.1162132.

[5] L. R. Rabiner, R. W. Schafer, and C. M. Rader, “The Chirp z-Transform Algorithm and Its Application,” *Bell System Technical Journal*, vol. 48, pp. 1249–1292, 1969. doi:10.1002/j.1538-7305.1969.tb04268.x.

[6] C.-P. Fan and G.-A. Su, “Pruning fast Fourier transform algorithm design using group-based method,” *Signal Processing*, vol. 87, no. 11, pp. 2781–2798, 2007. doi:10.1016/j.sigpro.2007.05.012.

[7] F. J. Harris, “On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978. doi:10.1109/PROC.1978.10837.

[8] FFTW Project, “Pruned FFTs with FFTW,” technical note, accessed 8 October 2026.

A. Reference Equations

$$
\begin{aligned}
X_F[m] &= \sum_{n=0}^{N-1} x[n]\mathrm{e}^{-j2\pi nm/N}, \\
X_E[m] &= \frac{N}{w}\sum_{n\in I_k} x[n]\mathrm{e}^{-j2\pi nm/N}, \quad m \in I_k, \\
E_w[m] &= \left(\frac{N}{w} - 1\right) S_k[m] - R_k[m], \\
D_{\mathrm{op},F}(w) &= \sqrt{\frac{N}{w} - 1}, \\
C_{\mathrm{direct}} &= Nw, \\
L(w) &= 2^{\lceil \log_2(2w-1)\rceil}, \\
C_{\mathrm{FastEDPW},\,principal} &= \frac{N}{w}[L\log_2 L + L], \\
C_{\mathrm{FastEDPW},\,opt} &\approx C_{\mathrm{FastEDPW},\,principal} + 2N, \\
C_{\mathrm{EDPW,best}} &= \min\{Nw, C_{\mathrm{FastEDPW},opt}\}, \\
C_{\mathrm{FastEDPW},\,sel,principal} &= B[L\log_2 L + L].
\end{aligned}
$$

B. Public Code Coverage and Verification Protocol

The independently packaged public release provides an automated experiment runner. Its first stage runs the MATLAB self-test, checking mathematical identities and algorithm agreement; an assertion failure aborts the experimental sequence. Subsequent drivers cover the wideband workload, frequency-grid accounting, and comparative spectral investigations, including the all-method common-bin and separate STFT tasks. A passing static dependency check or independent NumPy cross-check is *not* a passing MATLAB execution test. In particular, reported spectral-quality and detection metrics must be derived from actual recorded outputs, not from theoretical operation estimates.

B.1. Cold-call versus reusable-kernel arithmetic

Let $L = 2^{\lceil \log_2(2w-1)\rceil}$ and $K = N/w$. The principal full-spectrum Fast-EDPW convolution cost, under a radix-2 complex-butterfly model, is

$$C_{\mathrm{conv}} = K\,[L\log_2 L + L]. \tag{85}$$

When the chirp kernel FFT is *computed once per call and reused across blocks*, its additional butterfly cost is $(L/2)\log_2 L$. The MATLAB cost helper therefore distinguishes the warm-

kernel estimate

$$\widehat{C}_{\text{warm}} = C_{\text{conv}} + 2N \tag{86}$$

from the cold-call estimate

$$\widehat{C}_{\text{cold}} = \widehat{C}_{\text{warm}} + \frac{L}{2}\log_2 L. \tag{87}$$

These remain *estimates*, not exact instruction counts: the $2N$ term assumes optimized phase-factor preparation, whereas the distributed MATLAB helpers may calculate additional complex exponentials and intermediate products. For a selected single block, replace K by 1 and $2N$ by $2w$ in the corresponding estimate. The kernel setup cost must not be omitted when the supplied code computes it.

For a CZT producing M coefficients from N samples, the convolution FFT length is $L_c = 2^{\lceil \log_2(N+M-1)\rceil}$. The analogous warm estimate is $L_c \log_2 L_c + L_c + N + M$, with an additional $(L_c/2)\log_2 L_c$ for a cold-call chirp-kernel FFT. The radix-2 butterfly reference $\frac{N}{2}\log_2 N$ is exact under its stated counting convention only when N is a power of two; otherwise it is an extrapolated reference, not a measured mixed-radix FFT count.

C. Reproducibility Status

The archived baseline remains unchanged. The analysis extends the analytical comparison and introduces a MATLAB benchmark. Full quantitative claims require MATLAB R2016a execution, operation-count audit, and manuscript–code–figure reconciliation. This statement is intentional: a theoretical benchmark design must not be presented as completed numerical validation.