

Persistent Modular Phase-State Updates for GFSK Waveform Generation

A Reproducible WSJT-X Case Study

Rudolf Stepan

Independent researcher, Vienna, Austria

OE3SRO | ORCID: 0009-0004-2842-2579

15 October 2026 | Journal research draft, version 0.1

Not peer reviewed; requires human scientific and editorial revision.

Abstract

Maintaining a bounded modular phase state can avoid general argument reduction in recursive oscillators. We investigate this established design principle in the Gaussian frequency-shift-keying generators of an implemented WSJT-X fork. Stepan’s REIST-Division manuscript motivates the persistent centered-state formulation; this study evaluates its floating-point application rather than claiming a new modular arithmetic method. Controlled alternatives include a positive phase interval, source-level masked correction, general remainder, 48-bit fixed-point DDS, an unbounded accumulator and two block-prefix designs. The production implementation remains unchanged. On one Intel Core i9-14900K, three process sessions per verified core class show that architecture choice depends materially on scheduling: the general-remainder/centered median ratio is approximately 2.6 on the selected performance core and 4.4–5.2 on the selected efficiency core. Positive and centered bounded states have similar costs. The tested fixed-point and prefix implementations are slower, without establishing a general ranking of DDS architectures. Numerical tests cover valid encoded frames, long trajectories, boundary cases and software spectra. Separate LDPC experiments demonstrate conditional success–cost improvements over BP alone, while an adaptive Costas search loses valid peaks. These decoder findings are unrelated to REIST and do not establish improved radio sensitivity. Raw data, exact source hashes, automated analysis and explicit limitations accompany the draft.

Keywords: modular phase accumulation; GFSK; DDS; reproducible benchmarking; floating-point arithmetic; WSJT-X.

1 Introduction

An oscillator can represent phase by a growing angle or by a persistent modular state. Both are mathematically valid, but their arithmetic cost, rounding and integration requirements differ. In a narrowband GFSK generator, every sample depends on the previous phase and a smoothly varying frequency increment. Replacing a general remainder with a justified bounded update may shorten this dependency chain. It does not imply that transmitted signals become easier for another station to receive.

Stepan’s REIST-Division manuscript [1] formulates centered integer remainders as persistent states, with a bounded additive correction. Its Section 6 explicitly discusses phase windows. Our application uses the same state-management idea with binary floating-point phase and a fixed represented period. The elementary invariant is independently proved below. The original paper’s CPU and FPGA speedups are not imported as evidence for this application.

The historical German application report [2] compared centered correction with general remainder while controlling the trigonometric strategy. It did not benchmark positive states or fixed-point DDS and did not isolate performance and efficiency cores. The present extension asks:

1. Under what arithmetic assumptions does the bounded update retain a canonical phase state, and what errors remain?
2. How do alternative phase architectures compare under identical pulse, frame, output and compiler conditions on this heterogeneous processor?
3. How much of a waveform-kernel gain can matter to a complete file decoder?
4. Can separate synchronization and LDPC experiments improve success per measured computational cost without obscuring losses or false accepts?

The contribution is a source-audited implementation case study and a reproducible comparison, including negative results. We claim neither priority for modular oscillators nor universal superiority of centered arithmetic.

2 Related work and attribution

Conventional DDS uses an accumulated phase register and amplitude mapping; frequency resolution and phase truncation are familiar design considerations in established technical treatments [3]. A floating state in $[0, T)$ is therefore an essential comparator, not an inferior representation by definition. General remainder supports larger domains and a specified tie convention, whereas a bounded update exploits additional input invariants.

Floating-point arithmetic distinguishes exact identities from rounded implementations. Cancellation and exact-subtraction conditions matter to the bounded update [4]. Prefix scans are established parallel building blocks [5]; applying them to phase removes some recursive work but reassociates additions and introduces extra work and storage.

FT4 and FT8 combine GFSK, synchronization patterns and error correction as documented by their designers [6]. Layered LDPC schedules and min-sum approximations have established prior art; for example, Cui et al. [7] study column-layered hardware decoding. Our experimental row-layered software kernel is not their architecture and is not a new discovery of layered decoding.

Table 1: Prior methods and the actual relationship to this study.

Method/source	Principle	Difference and empirical role
Persistent centered state, Stepan [1]	Integer state plus bounded signed correction	Motivates the application. Floating invariants and GFSK tradeoffs are evaluated here; earlier benchmark numbers are not reused.
DDS/NCO [3]	Modular phase register, amplitude mapping	Positive floating state and Q48 integer state are conventional competing implementations.
General remainder	Reduce a broadly valid input each update	Same double increments and LUT; measured reference with broader functionality.
Prefix scan [5]	Reassociate partial sums	Serial and tree block-64 architectures; additional reductions and rounding are included.
Hybrid FT decoding [6]	Iterative BP and OSD	Baseline for a separate decoder study; no modular-phase or REIST contribution.

The supplied REIST PDF identifies Rudolf Stepan, the stated title, year 2026, ORCID and DOI [1]. Its SHA256 is recorded in the audit. Public DOI and Zenodo metadata could not

be retrieved in this investigation, so the DOI is explicitly an author-manuscript identifier, not independently verified metadata. An unrelated DOI in that manuscript’s reference [3] was found to identify a different publisher article. That reference is not inherited here. The application relationship rests on the supplied manuscript’s definitions and additive invariant, not on that erroneous citation.

3 Mathematical foundations

3.1 State invariant over real numbers

Let $T > 0$ and define the canonical centered representative

$$C_T(x) = x - T \left\lfloor x/T + \frac{1}{2} \right\rfloor \in [-T/2, T/2). \quad (1)$$

For $r \in [-T/2, T/2)$ and $u \in [-T/2, T/2]$, the sum $s = r + u$ lies in $[-T, T)$. Consequently

$$r^+ = \begin{cases} s - T & s \geq T/2, \\ s + T & s < -T/2, \\ s & \text{otherwise} \end{cases} \quad (2)$$

is congruent to s and lies in the canonical interval. The three cases map to $[-T/2, 0)$, $[0, T/2)$ and the unchanged interior, respectively. Induction establishes persistence. This is the additive-state application of the original REIST formulation [1, Theorem 3.6 and Corollary 3.7].

For a positive state $p \in [0, T)$ with the same signed step bound, $p + u$ lies in $[-T/2, 3T/2)$. One addition for a negative sum or one subtraction for a sum at least T also suffices. Thus the absence of general division follows from bounded state and increments, not uniquely from centering.

3.2 Binary floating-point invariant and error

Write $P = \text{double}(\pi)$ and $\hat{T} = 2P$, where multiplication by two is exact. Assume binary64, round-to-nearest, finite stored $r \in [-P, P)$, finite stored $u \in [-P, P]$, and no reassociation changing the specified operations. Monotone rounding implies $s = \text{fl}(r + u) \in [-2P, 2P]$; the rounded upper endpoint may be attained. If $s \geq P$, subtraction of $2P$ is exact by the factor-of-two exact-subtraction condition (Sterbenz’s lemma). If $s < -P$, adding $2P$ is the corresponding exact subtraction of comparable positive magnitudes. These cases yield $[-P, 0]$ and $[0, P)$; the interior already lies in $[-P, P)$. The binary state invariant therefore holds under these explicit conditions, including a rounded sum equal to $2P$. It is stronger than merely testing real arithmetic, but is not a proof for NaNs, changed rounding modes, fast-math transformations or arbitrary public API inputs.

For normal sums, the only local state-update rounding is the first addition, bounded absolutely by $\epsilon_u \hat{T}$, where $\epsilon_u = 2^{-53}$. With gradual underflow an additive subnormal term can be included. A conservative circular accumulation bound for n already represented steps is

$$d_{\hat{T}}(\hat{r}_n, r_n) \leq n\epsilon_u \hat{T}, \quad d_T(a, b) = \min_{k \in \mathbb{Z}} |a - b - kT|, \quad (3)$$

up to subnormal terms. This is a worst-case bound, not a drift prediction. It excludes errors in constructing the increments and in representing true 2π , and any error from the large-step remainder guard.

General `std::remainder` uses a nearest-integer tie rule and can return $+P$ at an exact half-period, unlike the canonical centered state. A simple positive update can also produce stored $\hat{T}$

when adding $\hat{T}$ to a tiny negative subnormal rounds upward. These states remain periodically meaningful; canonical interval membership and oscillator equivalence must not be confused. Dedicated boundary data record both cases.

3.3 Fixed point and reassociation

The experimental Q48 state represents cycles by $q/2^{48}$. Its update is

$$q^+ = (q + \text{round}(2^{48}u/\hat{T})) \bmod 2^{48}. \quad (4)$$

Unsigned addition and masking are exact. Under exact conversion, step quantization contributes at most $n\hat{T}/2^{49}$ to circular error; the actual double scaling and rounding must also be considered. Small steps can quantize to zero. This implementation converts each varying GFSK step with `llround`; it is not a claim about the best possible DDS conversion path.

In a block of $b = 64$ samples, offsets are prefix sums of the increments. They may span several turns, so one bounded correction is insufficient for each offset. The experiment uses a floor-based reduction for emitted offsets and remainder for the block carry. Serial prefix accumulation retains a short addition chain; a tree scan exposes independent work at additional arithmetic and memory cost. Both alter rounding order. A two-array tree workspace occupies 1024 bytes of doubles, in addition to the unchanged waveform working buffers.

4 Algorithmic design and WSJT-X implementation

The production generator uses double increments and a centered double phase, then linearly interpolates a 65,536-interval complex float table. The ideal linear interpolation bound $h^2/8$, $h = 2\pi/65536$, is about $1.15 \cdot 10^{-9}$ per component; float table storage and rounded interpolation are separate error sources. Phase is sampled before the update. Ramps and Gaussian-pulse construction remain mode-specific and unchanged.

Table 2: Frame parameters at 48 ksample/s in this fork.

	FT8	FT4	FT2
Modulation tones	8	4	4
Input symbols	79	103	103
Samples/symbol	7680	2304	1152
Output samples	606720	241920	120960
Duration (s)	12.64	5.04	2.52
Gaussian BT	2	1	1

The code is identified by commit `46ebbd78c663bf07eed054570df4888a6268e156` and new experiment-file hashes. Every one of the historical manifest’s 45 file hashes matches that commit. The historical description of an uncommitted snapshot was accurate at the time; the original manifest is preserved rather than rewritten retrospectively.

`journal/experiment` above abbreviates `docs/papers/reist-ftx/journal/experiment`. The original Fortran waveform sources remain available for reference tests. Their different precision, trigonometry and caching preclude a language-performance interpretation. FT2 uses the fork’s FT4-based integration; it is not an additional independent LDPC code. Native C++ experiments do not imply that remaining Fortran routines are interpreted: they already execute compiled machine code.

Table 3: Implementation locations; paths are relative to the repository root.

File/function	Role
<code>lib/ft8/native/ft8_wave.cpp</code> , <code>waveform</code>	Common FT8/FT4/FT2 phase, LUT, pulses, guards and ramps.
Same file, <code>pulseFor</code>	At most eight immutable shared pulse objects; one cache lock per frame; scratch belongs to each invocation.
<code>journal/experiment/phase.hpp</code> , <code>oscillate</code>	Isolated alternative oscillators; generated copies replace only the oscillator segment.
<code>lib/ft4/native/ft4_sync.cpp</code> , <code>search</code>	Reuses frequency-tweaked Costas references; independent rows may be parallel.
<code>lib/ft8/native/ft8_ldpc.cpp</code>	Production native encoder and two BP parity layouts; unchanged in this study.
<code>lib/ft8/decode174_91.f90</code>	Standard hybrid BP/OSD, also called by FT4/FT2.
<code>tests/ft8/research/layered.cpp</code>	Experimental row-layered SPA/min-sum; private per-call state, excluded from production.

5 Experimental methodology

5.1 Architecture experiment

The machine is an Intel Core i9-14900K, 24 physical cores and 32 logical CPUs, running Windows 11 Enterprise build 26300. The toolchain is MSYS2 GNU GCC/GFortran 15.2.0. Waveform experiments use C++11, `-O3 -fno-fast-math -ffp-contract=off`, without LTO or `-march=native`. A separate Boost 1.90 decimal-100 check needs C++14; that language-standard requirement does not alter the timed kernels.

Per-logical-CPU affinity and CPUID leaf 1A identify type 0x40 and 0x20 cores [8]. The selected performance thread is logical CPU 0 and the selected efficiency thread CPU 16. SMT siblings, Turbo, power management and background load are not isolated. Three additional sessions allow free scheduling. The results apply to these selected cores, not every P/E core.

Each core class has three process sessions, nine paired randomized blocks per mode and ten implementations. A block contains three complete frames per implementation after warm-up. The new timed output is complex float IQ. All variants receive the same encoded telemetry word (seed 991), 48 ksample/s and 1855 Hz. Telemetry bits are encoded with the existing LDPC/CRC and proper Costas, Gray and FT4 scrambling mappings. Pulse synthesis, allocation, cache, ramp work and output size are included. The refactored centered oscillator must match production output bytes on the numerical corpus.

The main estimator is a ratio of pooled medians. Conditional uncertainty intervals resample whole process sessions and paired circular blocks of length three within a session (2000 replicates, seed 991202610). Lengths one and five provide sensitivity checks. The resampling preserves local timing dependence; with only three sessions, coverage is not established, and shared thermal/OS effects can correlate even separate processes. These are conditional resampling intervals, not guaranteed population confidence intervals or cross-machine inference. Individual repeated frames are not treated as independent samples.

Wall time uses `steady_clock`. Windows `QueryThreadCycleTime` supplies accounted user-plus-kernel thread cycles, normalized per sample, as a diagnostic [9]. These are not retired instruction counts or a calibrated PMU measurement. Hardware branch misses, cache events, energy and dynamic instructions/sample were not collected. Compiler assembly is supplied for small update loops; source-level masking is not assumed to guarantee branch-free machine code in every

context.

5.2 Historical reproduction, numerical corpus and spectra

The historical real-output harness is rerun into new files using the original 21-block, ten-repeat, seed-991 tone-index workload. Its random valid tone indices need not encode a message. These data are kept distinct from the new complex-output architecture study and from the frozen original measurements.

The numerical architecture corpus has three modes, two sample rates (12/48 ksample/s), two encoded telemetry messages, and five carriers per rate: -1855 , 0.001 , 1855 , $0.49f_s$ and 10^9 Hz. Thus it contains 60 frame configurations and ten implementations. The large carrier exercises the same general-remainder step guard, not realistic RF operation. The long-double reference accumulates the same represented double increments modulo the same double period and evaluates long-double sine/cosine. Its 64-bit significand is an improved arithmetic reference, not an exact reference for Gaussian pulse construction or physical pi. Decimal-100 calculations independently check 4096 local updates; a million random state/step pairs check interval membership.

Separate million-step trajectories use a constant positive carrier step, signed random steps and tiny alternating steps. Periodic phase error, IQ maximum/RMS error and actual byte equality are distinct outputs. Software spectra use the complete ramped frame, rectangular observation, zero-padding to the next power of two and a deterministic radix-2 FFT. A known-bin self-check and Parseval error validate the transform. A chosen analysis band extends two tone spacings beyond the nominal lowest/highest tones; this is not a regulatory emission mask. Error energy is normalized to reference-frame energy. All quantities concern generated digital IQ.

5.3 Complete-decoder component profiling

An explicitly excluded executable includes the unchanged production waveform source and times its ABI calls inside the complete `jt9` file decoder. Recorded FT8, FT4 and optional local FT2 samples run at depth three, one OpenMP thread, three measured processes after one warm-up. Output message sets must remain identical across runs. The reported component share divides summed waveform time by process launch-to-exit wall time; it includes startup and is specific to these files. It is not a GUI, live-stream or transceiver profile.

5.4 Independent decoder and synchronization studies

The first decoder pilot uses 128 words/cell, seed 20261016. A separate confirmation uses 2048 words/cell and seed 20261017. Five scenarios comprise BPSK AWGN LLRs, a 12-bit zero-LLR burst, 28 correct frozen AP bits, one wrong frozen AP bit, and noise alone, each at $\sigma = 0.8, 1.0, 1.2$. The channel model is $y = b + \sigma z$, $b \in \{-1, 1\}$ and $L = 2y/\sigma^2$; positive LLR means bit one. Its $E_s/N_0 = 1/(2\sigma^2)$ is a BPSK channel coordinate, *not* FT-mode SNR in a 2500 Hz bandwidth. Payloads and noise differ between trials and match between methods. Box–Muller and explicit `mt19937`-to-uniform conversion avoid library-specific normal-distribution transforms.

Both parity-row layouts are tested. Baselines are native flooding BP30 and, for the standard layout, the existing hybrid with two saved BP reliability sums and OSD order two. Alternatives are row-layered sum-product (15 sweeps) and normalized min-sum (30 sweeps, $\alpha = 0.75$ or 0.85 , plus 60 sweeps at 0.75). SPA check messages are bounded at approximately ± 14 . The layered state updates immediately after each check row. AP values remain frozen.

Following the confirmation, two data-informed cheap-first portfolios are specified: MS085 then unchanged BP30, or MS085 then the standard hybrid on failure. They use a new 2048-word/cell

dataset, seed 20261018, without pooling it with confirmation data. An accepted wrong word could preempt fallback; success preservation is therefore an empirical result, not a universal claim. The SMP-specific OSD is not benchmarked by substituting the standard hybrid.

Every success must satisfy syndrome and CRC checks and match the known full 174-bit truth. Incorrect accepts, including zero words, are retained even when the application would apply additional semantic filters. Paired gains/losses are reported. Marginal Wilson intervals describe each correctness proportion; they are not intervals for paired differences. Times cover serial decoder calls in randomized 32-word batches, excluding data generation and CSV I/O. This is a mean-cost comparison, not a hard per-call deadline guarantee.

The Costas experiment injects four reference arrays in complex noise and compares the existing full grid (585 points) with an 85-point coarse grid plus a small local refinement. It uses 1024 trials per real/imaginary noise standard deviation 1, 2, 4, 8, seed 174202610, random time and frequency, and one thread. A localization counts as correct within two downsampled samples and one Hz. It measures a simplified search-kernel proxy, not complete frame decoding, FT8 search recall, interference cancellation or a measured RF SNR curve.

6 Performance results

Table 4: Complex-output waveform time on the selected P-core: median of 27 block means, milliseconds per complete frame. All pulse/ramp work included.

Architecture	FT8 (ms)	FT4 (ms)	FT2 (ms)
Centered	2.774	1.088	0.550
Remainder	7.201	2.857	1.431
Positive	2.860	1.089	0.601
Masked	2.876	1.134	0.557
Q48 DDS	6.482	2.672	1.332
Unbounded	10.857	4.347	2.227
Serial prefix-64	4.758	1.890	1.024
Tree prefix-64	5.864	2.403	1.260
Long double	27.437	10.588	5.386
Production	2.719	1.038	0.528

Table 5: Core dependence of general remainder (G) versus centered (C). Intervals are conditional hierarchical resampling intervals; see methodology.

Core	Mode	G/C	Conditional 95% interval	C (ns/sample)
P	FT8	2.60	[2.35, 2.71]	4.57
P	FT4	2.62	[2.32, 2.75]	4.50
P	FT2	2.60	[2.53, 2.97]	4.55
E	FT8	5.19	[5.05, 5.33]	5.96
E	FT4	4.94	[4.70, 5.08]	6.37
E	FT2	4.43	[4.27, 4.61]	7.04
UNBOUND	FT8	2.63	[2.62, 2.67]	4.16
UNBOUND	FT4	5.03	[2.41, 5.17]	6.46
UNBOUND	FT2	5.53	[1.98, 5.85]	5.62

Positive bounded states are close to centered states in these measurements. The practical saving is principally avoiding repeated general reduction. Source-level masking does not yield a consistent improvement. The tested Q48 conversion and block scans cost more than scalar

centered correction. These are negative results for the actual implementations, not proofs that optimized integer DDS or explicit SIMD cannot win.




Figure 1: FT8 architecture cost normalized separately within each core class. The figure shows point estimates; interval and sensitivity data are provided in the machine-readable summary.

Table 6: Separate real-output rerun of the historical factorial harness. These ratios should not be conflated with the new complex-output comparison.

Mode	Historical G-L/C-L	New real-output reproduction
FT8	5.55	5.56
FT4	5.24	5.21
FT2	4.99	5.19

The rerun tests repeatability under the current uncontrolled machine state, not exact equality of elapsed times. Differences between P/E classes and unbound processes show why the earlier 5.55/5.24/4.99 factors cannot be elevated to machine-independent constants. Compiler context and output type also differ between the historical and expanded harnesses.

Parallel experiments use independent centered waveform calls with 1/2/4/8 workers, five blocks and eight frames per worker. They include thread creation, scratch allocation and cache locks, and deliberately exclude core-bound sessions. The complete raw throughput table is retained; this is throughput of independent frames, not SIMD acceleration of one recursive oscillator or proof that a twelve-worker decoder scales similarly.

7 Numerical accuracy and signal validation

The centered refactor matches production bytes throughout this corpus. Other architectures need not be byte identical to be periodically equivalent. For a million constant positive steps, centered and general states have a maximum periodic error of approximately $6.86 \cdot 10^{-11}$ rad relative to the long-double represented-period reference. Positive accumulation reaches $1.53 \cdot 10^{-10}$ rad, Q48 about $2.38 \cdot 10^{-9}$ rad and the unbounded accumulator about $4.40 \cdot 10^{-6}$ rad. Signed random steps behave differently; these figures must not be interpreted as universal drift rates. Tiny positive steps can disappear entirely in Q48; boundary tests exhibit the general-remainder tie and positive-state endpoint effects discussed earlier.

Table 7: Worst result across the 60-frame corpus versus the long-double reference. IQ maximum and RMS use complex magnitude; phase excludes ramps whose reference amplitude is at most 0.1.

Architecture	Maximum IQ error	Maximum RMS	Maximum phase error (rad)
Centered	8.43e-08	3.20e-08	1.30e-07
Remainder	8.43e-08	3.20e-08	1.30e-07
Positive	8.43e-08	3.20e-08	1.30e-07
Masked	8.43e-08	3.20e-08	1.30e-07
Q48 DDS	8.43e-08	3.20e-08	1.30e-07
Unbounded	2.22e-06	8.80e-07	2.27e-06
Serial prefix-64	8.43e-08	3.20e-08	1.30e-07
Tree prefix-64	8.43e-08	3.20e-08	1.30e-07




Figure 2: One generated FT8 telemetry frame at 1855 Hz and 48 ksample/s. Error spectra use IQ differences, not the difference between logarithmic power curves. This is a software comparison, not a transmitter measurement.

Integrated centered/reference error energy is about -150 dBc for the three selected frames; the Q48 implementation is similar at this float-output floor. Parseval relative errors are below $7 \cdot 10^{-13}$. Reference energy outside the explicitly chosen analysis band is approximately -38.1 , -48.3 and -50.9 dBc for FT8/FT4/FT2. Those values reflect the observation, ramp and chosen band and are not spectral-mask certifications. They do not demonstrate reduced RF sidelobes. The long-double reference still shares represented pulse increments, final float conversion and ramp coefficients with the variants.

8 Practical relevance and independent decoder results

Table 8: Unchanged waveform component inside the complete single-thread recording decoder; medians of three runs after warm-up. The ideal bound removes this component entirely while holding all other work constant.

Mode	Wall (ms)	Wave (ms)	Calls	Share (%)	Ideal bound
FT8	1574.7	33.47	254	2.13	1.022
FT4	177.7	6.63	18	3.73	1.039
FT2	144.6	3.46	5	2.39	1.024

For baseline component share p and component speedup S , Amdahl’s model is

$$S_{\text{total}} = \frac{1}{1 - p + p/S}, \quad S_{\text{total}} \leq \frac{1}{1 - p}. \quad (5)$$

Here the measured share belongs to the current centered implementation. To model replacing it by a slower general-remainder implementation with relative cost G , the corresponding counterfactual ratio is $1 + (G - 1)p$. This is a conditional model; mode-dependent sample rates, cache state and input mix can change G . We do not substitute a 48 ksample/s TX ratio for an exact measured RX ratio. The profile establishes where the current component matters, not a measured full-program improvement from a hypothetical replacement.

Table 9: Separate portfolio dataset: standard parity layout, BPSK AWGN, 2048 independent words per noise level. Time is mean batch cost relative to BP30. False accepts count accepted words different from the known truth.

Noise σ	Method	Correct/2048	Gain/loss	False	Time/BP
0.8	BP30	1377	0/0	0	1.00
0.8	Hybrid	1971	594/0	0	1.75
0.8	MS085	1499	132/10	0	0.13
0.8	MS→BP	1509	132/0	0	0.63
0.8	MS→Hybrid	1971	594/0	0	1.22
1.0	BP30	56	0/0	0	1.00
1.0	Hybrid	587	531/0	0	3.81
1.0	MS085	98	42/0	0	0.31
1.0	MS→BP	98	42/0	0	1.28
1.0	MS→Hybrid	589	533/0	0	4.02
1.2	BP30	0	0/0	0	1.00
1.2	Hybrid	31	31/0	0	6.61
1.2	MS085	3	3/0	0	0.59
1.2	MS→BP	3	3/0	0	1.58
1.2	MS→Hybrid	31	31/0	0	7.34

At $\sigma = 0.8$, the new portfolio dataset has 1377 BP successes, 1499 MS085 successes and 1509 MS-to-BP successes. MS085 gains 132 cases and loses 10; the portfolio retains those ten via

fallback. Its measured average cost is about 0.63 of BP alone. This answers the limited kernel question affirmatively for this distribution: higher correct-word probability can coexist with a lower measured mean budget. It is not a hard-deadline or full-receiver result. At $\sigma = 1.0$ and 1.2, the same portfolio costs more than BP because most frames reach fallback. There is no uniform dominance across channel conditions.

The standard hybrid remains substantially more successful than BP/min-sum alone. A min-sum prepass occasionally helps or reduces its cost, but also adds overhead in weak/noise cases. The complete confirmation and portfolio matrices retain AP, wrong-AP and noise controls, both parity orders, and all tested variants. A CRC-valid false word occurs in the original hybrid on the confirmation AWGN $\sigma = 1.2$ dataset; this is a kernel false accept before application semantic filters, not a newly introduced on-air false message. Zero observations in other finite cells do not prove a zero false-decode rate.




Figure 3: Three sampled channel levels in the independent portfolio dataset. Lines guide the eye; they are not fitted sensitivity curves. Wilson intervals are available in the accompanying CSV and are marginal, not paired tests.

Table 10: Simplified FT4 Costas localization: 1024 paired trials per noise level. Time tolerance is two downsampled samples, frequency tolerance one Hz.

Noise σ	Search	Correct/1024	Gain/loss	Median (μ s)
1	full	1024	0/0	122.4
1	adaptive	1022	0/2	26.9
2	full	1023	0/0	121.8
2	adaptive	977	0/46	25.8
4	full	524	0/0	122.5
4	adaptive	440	46/130	27.0
8	full	57	0/0	122.5
8	adaptive	53	21/25	26.8

The adaptive search is roughly four to five times faster but loses peaks even at the strongest tested level. At noise standard deviation four, it finds 440 correct localizations versus 524 for the full grid, with 46 gains and 130 losses. This is evidence against deploying this particular pruning rule as a sensitivity improvement. A production strategy would need several retained coarse peaks or a justified uncertainty-based fallback, followed by full-message tests. No new decoder algorithm from this study is enabled in the application.

9 Discussion

The comparison supports persistent bounded state as a useful implementation choice for this GFSK kernel. It does not establish a special arithmetic advantage of centering over every positive representation. Centering provides a convenient signed error coordinate and, under the stated binary64 conditions, exact corrective subtractions. The magnitude of the performance advantage against library remainder is strongly core- and context-dependent.

Q48 DDS trades state-rounding error for step quantization and conversion cost. The evaluated per-step `llround` implementation is not an optimized fixed-point standard to beat. Likewise, a work-increasing tree scan compiled for the default ISA is not an assessment of all SIMD implementations. Explicit AVX2 scans, fused pulse/phase representations, vector conversion and independent oscillator batching remain plausible subsequent studies. Reassociation would need its own error evidence; it should not be hidden by fast-math flags.

The separate decoder results illustrate why throughput and sensitivity must be measured together. A cheap prepass can reduce later work on favorable inputs, yet consume additional time when almost everything fails. A search that prunes most hypotheses may be faster while missing real peaks. Neither phenomenon is derived from REIST or from the transmit phase invariant.

10 Threats to validity

Only one physical machine and one compiler version are measured. No AMD or ARM64 result, hardware PMU event, RF measurement or energy result is supplied. P/E identity is verified but core isolation and power control are absent. Three process sessions and short blocks give limited information about thermal and scheduler dependence; conditional bootstrap intervals can understate systematic uncertainty. Ratios of pooled medians differ from medians of paired ratios. Repeated frames within a block are not independent experimental units.

The expanded oscillator harness uses complex output, different encoded inputs and a refactored compile context relative to the historical real-output factorial experiment. The historical rerun is therefore reported separately. Both include unchanged pulse synthesis and ramps, so the measured ratio is whole-kernel cost rather than the cost of an isolated addition.

Long-double references share double increments and period; decimal-100 local checks do not validate the complete Gaussian pulse to 100 digits. The numerical corpus is finite, includes only two telemetry messages per setting, and cannot prove universal bit identity. Spectra cover selected ideal generated frames; no audio interface, USB path, amplifier, antenna or FT-991A was measured.

LDPC BPSK LLRs omit FSK demodulation, synchronization errors, fading diversity, co-channel interference and application semantic rejection. Random 77-bit LDPC payloads are code-level messages, not all valid displayed message formats. The SMP-specific OSD strategy is not equivalent to the standard hybrid baseline. The Costas injection is a simplified decimated-reference model with no Gaussian transition or full decoder. Decoder probabilities cannot be translated into FT8/FT4/FT2 RF sensitivity or a dB improvement. Mean elapsed budget does not establish per-frame deadline behavior or total multi-station throughput.

There has been no external replication, independent mathematical review or journal peer review. This draft was substantially prepared with AI assistance. Its existence does not establish compliance with a particular journal’s authorship or AI policy. The human author must critically revise and take responsibility for any eventual submission.

11 Conclusions

Maintaining a bounded modular phase avoids unnecessary general reduction in the tested GFSK implementation. A precise arithmetic invariant explains when one correction is sufficient; actual error and performance still depend on representation, rounding, output mapping and architecture. The expanded comparison finds positive bounded state competitive with centered state and shows material P/E-core effects. It also documents slower alternative implementations and input-dependent decoder tradeoffs rather than selecting only favorable results.

The relationship to REIST is an application of its persistent centered-state formulation, supported by the supplied manuscript and independent derivation. The contribution beyond the earlier application report is a broader controlled comparison, arithmetic boundary analysis, core-specific measurements and separate success-cost decoder experiments. No broader mathematical priority, RF improvement or publication readiness is claimed.

12 Data and code availability

The application base is <https://github.com/rudolfstepan/wsjtx>. The source commit is 46ebbd78c663bf07eed054570df4888a6268e156. New uncommitted experiment files are identified by an immutable source manifest in the accompanying reproduction package. Frozen historical raw data remain in `docs/papers/reist-ftx/data`; new data are in the separate `docs/papers/reist-ftx/journal/results` tree. Generated tables and plots are derived by versioned Python scripts. Original and new datasets are never pooled silently. The supplied original REIST PDF and optional third-party FT2 recordings are not redistributed.

`Build-Study.ps1` is the documented entry point for builds, new measurement directories, analysis and LaTeX. Standalone waveform experiments need GNU C++, Boost headers, CMake, Ninja and Python; historical reproduction additionally needs Fortran. Complete decoder/profile tests need the full fork and its normal dependencies. The source code retains GPLv3 licensing and upstream credits. The manuscript is a draft; its eventual publication license is undecided. No new release, DOI registration, remote upload, commit or push is made as part of this study. A versioned archive and verified DOI can be prepared later with the author's explicit approval.

Author contributions and AI assistance

Rudolf Stepan supplied the REIST framework, existing source fork, research objectives and original manuscript. OpenAI Codex assisted with source audit, experimental implementation, local execution, analysis, diagrams and substantial draft text. Human verification and author-written revision remain required. This disclosure is not a substitute for journal-specific policy compliance. The work acknowledges the original WSJT-X protocol and software developers and the fork's documented FT2 contributors; no authorship or license credit is transferred from them to the phase-state application.

References

- [1] R. Stepan, *REIST-Division: Eine implementierungsorientierte Formulierung zentrierter Restarithmetik. Theorie, Architekturanalyse und empirische Evaluation auf ARMv8-A, x86-64 und FPGA*, supplied manuscript, 2026. DOI as printed: 10.5281/zenodo.21206471; public metadata not independently retrieved. Original PDF SHA256 in the audit.

- [2] R. Stepan, *REIST Division als Zustandsprinzip in den GFSK Rechenkernen von FT8 FT4 und FT2*, technical report v1.0, 15 October 2026, docs/papers/reist-ftx/paper.tex, archived with commit 46ebbd78c.
- [3] Analog Devices, *MT-085: Fundamentals of Direct Digital Synthesis (DDS)*, technical tutorial. <https://www.analog.com/media/en/training-seminars/tutorials/mt-085.pdf>.
- [4] D. Goldberg, “What Every Computer Scientist Should Know About Floating-Point Arithmetic,” *ACM Computing Surveys*, vol.23, no.1, pp.5–48, 1991. Authoritative reprint: https://docs.oracle.com/cd/E19957-01/806-3568/ncg_goldberg.html.
- [5] G. E. Blelloch, *Prefix Sums and Their Applications*, Carnegie Mellon University, CMU-CS-90-190, November 1990. <https://www.cs.cmu.edu/afs/cs.cmu.edu/project/scandal/public/papers/CMU-CS-90-190.html>.
- [6] S. Franke, B. Somerville and J. Taylor, “The FT4 and FT8 Communication Protocols,” *QEX*, July/August 2020, pp.7–17. https://wsjt.sourceforge.io/FT4_FT8_QEX.pdf.
- [7] Z. Cui, Z. Wang and X. Zhang, *Reduced-Complexity Column-Layered Decoding and Implementation for LDPC Codes*, arXiv:1204.2577, 2012. <https://arxiv.org/abs/1204.2577>.
- [8] Intel, *Intel 64 and IA-32 Architectures Software Developer’s Manual*, CPUID leaf 1AH, hybrid core-type enumeration. <https://cdrdv2-public.intel.com/868137/325462-089-sdm-vol-1-2abcd-3abcd-4.pdf>.
- [9] Microsoft, *QueryThreadCycleTime function*, Win32 API documentation. <https://learn.microsoft.com/en-us/windows/win32/api/realtimeapiset/nf-realtimeapiset-querythreadcycletime>.