

Passive Precision Time over BLE: Prospective Edge Calibration, Causal Publication, and Service Availability

Jonathan Yantis (Engineer)

Research preprint for local review — 2 October 2026

Abstract

Precision Time over BLE (PToB) investigates passive time distribution across heterogeneous receivers. We distinguish physical reference-edge delay, causal delivered-clock error and service availability. Separate bidirectional ADDRESS/PPS training froze a single 10,516 ns input correction before prospective validation on two nRF52840 clocks. Forward and reverse PPS-shadow measurements gave uncentered RMS 70.197 and 76.359 ns, MAE 57.263 and 60.376 ns, and absolute P99 158 and 219 ns on 600/600 and 598/599 physical PPS observations. Neither receiver's GPS trained its BLE clock; no validation refit, per-channel correction, centering or outlier removal was used. These MAE, SD and P99 values are numerically below MicroSync's published figures, with different metrology and no independent GPIO or energy comparison. An ESP32-S3 receiver separately achieved 319.5 ns internal causal RMS on all 359 physical targets in a normal 600 s run, while 1123/1198 sampled views exposed a fine published clock. Serving-aware authority selection reduced observed nonserving fallback episodes to 2.000–3.367 s; cold-start error and remaining publication gaps preclude full-service acceptance. The findings support profile-specific edge calibration and explicit causal/availability contracts, while delimiting output, holdover and power qualification still required.

1 Introduction

A passive time service lets a small clock, sensor or event recorder obtain useful time without transmitting, pairing, enrollment or a key. Different receivers need different precision and energy profiles. PToB separates public coarse UTC, lightweight timestamp follow-up and physical-stamp refinement from authorized diagnostics, elections and eventual mesh roles. Legacy advertising on channels 37, 38 and 39 is the selected transport.

The development question is broader than how closely a regression fits timestamps. A stable physical-edge delay can dominate delivered error while leaving small residual scatter. A good clock may remain unusable because the selected provider is not publishing. Conversely, recovering authority does not guarantee that the next physical original and its correction have completed. These effects need distinct measurements.

This paper contributes (1) a prospectively validated, common bidirectional edge-correction profile; (2) a causal measurement framework separating PPS-shadow, frozen hardware compare, event prediction and sampled availability; and (3) evidence that serving-aware authority protection shortens fallback without removing all service gaps. Exact physical-copy identity and epoch provenance connect these mechanisms. The work is not an adopted IEEE/SIG standard or a PTP BMCA/profile implementation; controller qualification remains separate from demonstrated timing behavior.

2 Related work and comparison definitions

FTSP combines low-level timestamps and skew regression in broadcast sensor networks [1]. CheepSync supplies prior-transmission timing information to heterogeneous passive BLE listeners, with reported accuracy in the $10\mu\text{s}$ range [4]. BlueSync uses connectionless BLE, PPI event capture and subsequent-packet timestamps; the author’s dissertation includes physical output comparisons [2]. Passive reception, follow-up and submicrosecond BLE timing therefore have clear precedents.

MicroSync uses connected nRF52840 devices, one central and three peripherals, PA/LNA timing signals and bidirectional exchanges. Its high-accuracy results are MAE 272 ns, reported SD 319.2 ns and 99% CDF 591.25 ns; MCU-only power is $252.1\mu\text{W}$. Low-power MAE is 796.4 ns at $141.8\mu\text{W}$. The evaluation uses external 800 MHz GPIO comparison and PPK2 excluding LEDs/peripherals. Equation (4) includes a premeasured, hard-coded delay d [3]. Thus an independently trained fixed edge correction is a legitimate experimental profile, provided its training and validation populations remain separate.

MicroSync does not report RMS, and we do not substitute its SD for RMS. Its reported SD is retained without assuming whether every plotted/aggregated population is identical to ours. Our signed-error population SD and absolute-error P99 are explicitly defined below. Numerical comparisons of the reported statistics are useful, but differing reference paths, topology, observation populations and instruments prevent a matched output-accuracy or energy claim.

Nordic’s reference design illustrates chain-delay accounting and physical output measurement using proprietary timeslot packets [5]. Zhou et al.’s 2026 preprint also uses proprietary radio timeslots: its approximately 20 ns result is at 1000 Hz without communication-stack activity, while BLE-active results are a separate traffic condition [11]. Those experiments do not measure passive legacy-advertisement timing under our profile. They reinforce the need to name the transport category instead of asserting an all-time BLE record.

2.1 Measurement classes

An event-fit residual describes a model at observed radio events. A causal prediction uses only prior available information. A same-copy receiver difference cancels shared source terms but not independent receiver error. A PPS shadow queries the already-published clock at a separately captured reference edge. A frozen TIMER/PPI compare evaluates a scheduled hardware event against adjacent PPS on a shared local counter. An external GPIO measurement includes the delivered electrical output and its instrument path. Absolute UTC additionally requires source calibration.

We retain these classes throughout. Internal hardware observations are stronger than a retrospectively centered fit, yet do not include an unconnected GPIO pad. Relative accuracy, UTC semantics, source quality and receiver delivered error remain separately reported. IEEE timing terminology is used for interpretation, not to assert profile conformance [9].

3 Timing model, physical identity and calibration

For local counter x , expressed in seconds after subtracting an integer anchor, a receiver may use

$$\hat{T}_{s,c}(x) = E + a_{s,c} + b(x - x_0) + \frac{1}{2}k(x - x_0)^2, \quad (1)$$

where E is a validated epoch, $a_{s,c}$ a source/channel intercept and b, k shared oscillator-rate terms. Source discontinuities require new contexts. Separate intercepts avoid interpreting channel bias as oscillator curvature, but do not identify absolute delay. T096 uses its existing BLE estimator; W12 evaluates bounded affine/curve alternatives with causal promotion. The model family does not imply identical estimators across platforms.

Each original has an immutable publisher/session, sequence, physical channel, PHY and RX capture. Correction metadata identifies that original even if received on another channel or scan epoch. Any two of seconds, fraction and XOR parity reconstruct the missing page; repeated pages add no independent observations. Conflicts, expiry, retired sessions and late contradictions cannot refresh a clock. A corrected historical timestamp is advanced from its original receive tick to the query instant, with explicit epoch/timescale and age. Small relative residuals alone cannot establish the correct UTC second.

3.1 A fixed physical-edge profile

On the nRF52840 path, both TX and RX are labeled ADDRESS events, but receive decoding introduces a reference-edge delay. Nordic’s hardware discussion confirms that RX ADDRESS follows decoding and can depend on signal quality [12]. Our correction is measured for the experimental profile rather than copied from a theoretical radio-delay value.

Separate training arms joined exact source UTC/session/sequence/channel to 361 forward and 265 reverse physical ADDRESS observations. Mean RX-minus-TX differences were 10,529.839 and 10,502.487 ns. The rounded directional half-sum,

$$d = \text{round}[(10529.839 + 10502.487)/2] = 10516 \text{ ns}, \quad (2)$$

was frozen before prospective validation. A stable pair-PPS phase term cancels in this half-sum under reciprocal/stable-path assumptions. Asymmetry, signal-dependent latency and later environmental changes are not thereby calibrated.

The deployed input rule is $M(x_{\text{RX}}) = T_{\text{TX}} + d$. The same d is added in both directions on every channel, with unchanged physical stamps and an identical-input parallel uncorrected model. It changes future BLE clock publication, not the recorded validation errors. No per-channel adjustment, validation mean subtraction or scoring-data fit is applied. Receiver GPS/PPS is shadow only and never trains BLE phase, rate or epoch.

This profile demonstrates compensation of a repeatable edge offset in the measured conditions. It does not transfer automatically to the ESP32-S3 descriptor path, different PHYs, boards, signal levels or relays. Root time origin and transmitted UTC semantics remain unchanged.

4 System and experimental methodology

T096 clocks TC01 and TC03 use nRF52840 hardware event capture on a 16 MHz counter (62.5 ns ticks). Either acts as the GNSS-root source while the other follows BLE. W12 C1/C2 use ESP32-S3, a 20 MHz local timer and a 500 ns receive-descriptor grid with controller-to-timer mapping. W12 has no validated independent UTC adapter. T096 has no persistent RTC adapter. Capture capability and source capability are therefore distinct.

4.1 Prospective T096 validation

Training was sealed and the 10,516 ns constant frozen before the two clean approximately 600 s validation arms. Source roles, image, transport and correction remained fixed; reversal used the same constant. A third reverse arm with returning field authority is retained as a failed service population. Actual/exported PPS opportunities, missing joins, ages, observer work and source epochs are reported rather than restricting results to a favorable subset.

The PPS shadow captures local PPS on the same hardware counter and queries the already-published BLE clock before local source processing. A publication completed after that edge cannot answer the historical query: such an observation is unscorable. The internal compare freezes the delivered model at arm time, schedules UTC+0.25 s with TIMER/PPI and compares its captured event to adjacent PPS. Every captured compare is one 62.5 ns tick after the planned event; that physical capture delay and grid rounding remain in the reported error. External GPIO is unconnected.

For signed error $e_i = \hat{T}_i - T_i$, we report bias $\bar{e}$, MAE $n^{-1} \sum |e_i|$, uncentered RMS $\sqrt{n^{-1} \sum e_i^2}$ and population SD $\sqrt{n^{-1} \sum (e_i - \bar{e})^2}$. P95/P99 are absolute-error ranks. No retrospective centering or error-dependent deletion is used. Source PPS 25 ns is the authorized working assumption, not an independently recertified absolute UTC bound.

4.2 W12 causal publication and authority

W12 predictions are frozen before each target trains the model. Model availability and fine/coarse/unavailable view counts use separate denominators. The final C2-only experiment consists of app-boot 90 s and normal 600 s captures separated by 82.746 s without observation; the gap is not continuous-service evidence. Earlier simultaneous C1/C2 runs remain a distinct profile.

Valid-time admission requires an advancing epoch, at least 60 s quiet boot and source/mesh stabilization; provisional RTC service additionally requires healthy isolated observation. Force, priority, MAC and ID cannot bypass admission. Mesh holdover cannot become an independent root. The selected policy reserves 10 s dwell/30 s tenure for a live qualified actual SERVING incumbent; a nonserving selected candidate uses 1.5 s/2 s. Claims, actual publication and UI share admission. Only the incumbent-protection condition changed in the final C2 profile, not wire bytes, authentication, lease, rank or estimator.

The final evidence cutoff is 17:28:13.400616 UTC on 2 October 2026. Exact images, seals, method hashes and all earlier populations are preserved in the companion manifest [13]. No hardware operation was performed for this manuscript update.

5 Results: prospective physical-edge compensation

Table 1: All scoreable corrected PPS-shadow errors, ns. Denominators are actual/exported PPS opportunities; no centering or tail deletion.

Quantity	Forward	Clean reverse	Interrupted reverse
Matched / all PPS	600 / 600	598 / 599	594 / 600
Signed bias	−11.137	+23.992	−66.394
MAE	57.263	60.376	108.875
Uncentered RMS	70.197	76.359	420.267
Population SD	69.308	72.493	414.989
Absolute P95	132	153	170
Absolute P99	158	219	1726
Absolute maximum	192	257	6011

The two clean directions confirm the fixed-correction mechanism under the stated warm profile. Parallel raw PPS clocks retain signed bias $-10527.137 / -10492.008$ ns and RMS $10527.365/10492.259$ ns. Their SD is unchanged by the constant correction. The improvement is therefore a prospective reduction of delivered signed offset, not a narrower post-hoc centered plot.

Clean reverse retains one physical PPS whose newer publication completed after its edge and before processing. It remains unscoreable. In the interrupted reverse, four such causal query gaps and two unavailable service observations remain among all 600 PPS. Returning field authority interrupted the serving source; the 594-scoreable distribution includes its recovery tails. Its low MAE alone cannot establish robust service.

Table 2: Frozen internal TIMER/PPI compare against adjacent PPS, ns. These are not external GPIO measurements.

Direction	Joins / captures	RMS	P99	Maximum
Forward	599 / 600	127.699	250	313
Clean reverse	598 / 600	102.237	234	250

The forward final join is end-censored. Clean reverse has one invalid reference join and one end-censored join. The one-tick capture delay and grid term are retained. During the clean benchmarks there were no missed events, pending overflows or benchmark log-drop increments; post-benchmark verbose drops remain in provenance. Maximum observer work was $6.010/5.976$ ms, which is measured work and cannot be declared negligible or converted into CPU percentage.

Both clean PToB matched populations have numerically lower MAE ($57.263/60.376$ versus 272 ns), population SD ($69.308/72.493$ versus reported 319.2 ns), and absolute P99 ($158/219$ versus reported 591.25 ns) than MicroSync [3]. This is a comparison of reported figures under named methods. MicroSync’s external GPIO and our PPS-shadow/shared-counter paths differ; MicroSync reports no RMS. The result establishes neither matched independent-output superiority nor an energy advantage, field certification or a BLE timing record.

6 Results: heterogeneous reception and recovery

Table 3 reports the latest C2 profile with all sampled views and physical targets. App-boot is a genuine application restart, not a characterized power-cold or temperature test. Prediction validity and render-state availability are different observations.

Table 3: C2 internal published prediction, with cold and normal populations separate. Errors in ns.

Quantity	App-boot 90 s	Normal 600 s
Scoreable / physical targets	47 / 57	359 / 359
Signed mean	+1342.9	+148.7
MAE	1348.4	244.5
RMS	1490.8	319.5
Absolute P95 / P99	2326 / 2474	666 / 930
Absolute maximum	2474	1010
Fine / coarse / unavailable views	144 / 21 / 15	1123 / 18 / 57
All sampled views	180	1198
Root changes	0	8
Longest unavailable interval (s)	7.524	11.503

All 359 normal targets were scoreable and RMS was below 400 ns. This confirms conditional prediction performance in the normal run. Cold error, the 1010 ns maximum and 57 unavailable views prevent full-service acceptance. Three inherited valid views have unverified prefix mappings; the truncated final view is retained. Fine publication on 1123/1198 views is useful sampled coverage, not continuous instrumented output qualification.




Figure 1: Latest C2 sampled view composition, exact counts in Table 3.

Four nonserving fallback episodes lasted 2.000, 2.002, 3.367 and 2.000 s, compared with approximately 30 s under the preceding protection rule. Canonical host cases isolate the timer semantics; the live observations demonstrate shorter selection recovery. The 11.503 s unavailable interval continues beyond authority return, separating source selection from physical-original/correction/publication recovery.

The preceding dual-receiver normal run remains distinct: C1/C2 scored 312/315 and 353/353 targets, RMS 243.2/248.2 ns, with fine views 902/1198 and 1066/1198 and longest blanks 30.013/30.008 s. Different reset, model and RF histories prohibit pooling or an identical-traffic A/B precision claim. The latest C2 same-input curve shadow has RMS 223.9 ns/max 597 ns; that is not the actual published 319.5 ns/max 1010 ns result.

7 What the experiments establish

7.1 Calibration must precede the claim

The bidirectional training/validation split converts a plausible fixed-delay hypothesis into a demonstrated profile mechanism. The uncorrected and corrected clocks use identical input observations, so the signed-offset reduction is directly attributable to the frozen input correction. The unchanged spread, common value across channels/directions and retained interrupted arm constrain that interpretation. They do not establish temperature, signal-strength or board-to-board invariance.

Physical capture, clock-domain mapping and fixed RF/path bias remain distinct. A 62.5 ns timer tick, a 500 ns descriptor grid and an externally observed GPIO edge provide different information. For illustration, uniform quantization over width q has RMS $q/\sqrt{12}$; a 500 ns grid gives about 144 ns. That is not a measured floor. Combining terms in quadrature additionally requires independence, which shared clocks and deterministic quantizer phase may violate.

7.2 Identity and provenance support both calibration and UTC

Exact same-copy joins establish that a correction refers to the measured physical original. Metadata carrier channel cannot overwrite that original's channel. Separate epoch/provenance checks prevent a free intercept from hiding a wrong UTC second. Same-copy receiver comparisons isolate local model effects by canceling common source terms, but do not establish either receiver's absolute UTC. Own GPS must remain outside follower training for the reference comparison to be informative.

Any-two metadata recovery handles one page erasure without turning duplicate reports into extra observations. Expiry, conflicts and late contradictions are part of the clock contract. These invariants allow reception on inexpensive clients while giving capable receivers a defensible refinement path.

7.3 Selection and publication are distinct service stages

The C2 recovery result establishes that long incumbent protection should track evidence of actual service, not candidate quality alone. The change is narrow: a qualified nonpublishing fallback no longer earns the serving incumbent's long dwell/tenure. Authentication, admission and source ranking retain their independent roles. Shorter fallback does not repair every original/correction or estimator gap, as the longer unavailable-view interval demonstrates.

A bounded sample count is not a bounded history age. Shared crystal rate/curvature with source/channel intercepts still needs fresh data, span, rate limits and causal promotion. The W12 learner retains minimum-six observations in 60 s, cell/span safeguards and the 100-pair no-tail-regression selector. A better shadow curve cannot be retroactively substituted for the published clock. Hardware scheduling and monotonic observation-time consistency likewise determine whether an otherwise good estimate is available at the requested instant.

7.4 Holdover is its own experiment

Earlier frozen-output evidence commanded 20 s without BLE but produced about 24 s between originals, including age and recovery. First internal forecasts were $-3183/-3009$ ns. This failed the future ± 500 ns total gate and is not superseded by clean warm timing runs. A held clock must be scored throughout its actual unavailable-reference interval; endpoint residuals or recent percentiles cannot provide that bound.

8 Limitations and next discriminating experiments

The clean T096 results establish warm prospective compensation in two directions. They do not prove corrected-profile reboot, source loss, cold acquisition, holdover or temperature robustness. The interrupted reverse and W12 cold populations show why these conditions must remain explicit. Single runs, correlated samples and different metrology prevent a general performance ranking. Repeated matched trials should use predefined phases and block-aware statistical analysis.

The PPS shadow is independent of BLE training, but it shares the receiver’s counter and an assumed GNSS reference. Pairwise GNSS phase/sawtooth, path asymmetry and absolute UTC are not independently calibrated. Source PPS 25 ns remains the authorized working assumption. External GPIO is unconnected; delivered electrical output therefore remains unmeasured. A regular independent output grid, including unavailable intervals, calibrated cables, instrument uncertainty and documented reference edges is the next accuracy test. Timestamp resolution alone is not an uncertainty certificate.

8.1 Resources and power–precision profiles

The measured T096 application/static footprint is 133,513/43,761 bytes. Final C2 is 757,808/177,716 bytes with minimum free heap 107,796 and main-stack watermark 7,144 bytes. C2 background-work maxima are 2771 μ s cold and 3121 μ s normal; T096 benchmark observer maxima approach 6 ms. These are size, remaining capacity and observed routine costs, not CPU busy percentage or a negligible-observer claim.

CPU percentage, actual listening duty/airtime, whole-device power/energy and battery life remain unmeasured. A useful power model combines mutually exclusive operating-state occupancies with transition energy per unit time, avoiding double counting concurrent CPU/radio states. The prospective 10 μ s client target is a desired maximum error, not a precision floor and not the 10,516 ns edge correction. Compare 1 Hz, 0.5 Hz and slower passive observation with acquisition, recovery, CPU/RX caps and work/residence costs measured explicitly.

Future BLE/PTS requires per-unit TOTAL held delivered-time error within ± 500 ns throughout actual blackout plus age, switching and synchronization margins, and at most 50% actual BLE listening. PTS remains off. The planning relation $U(h) \leq U_0 + U_f h + \frac{1}{2} U_a h^2 + U_{\text{switch}} + U_{\text{path}}$ becomes a bound only after its terms are qualified. At least 100 distinct frozen blackout trials per unit over a 30-minute warm run are the proposed minimum; repeated 30-minute profile trials and a 24-hour thermal/coexistence soak remain planned.

8.2 Adapters and forwarding

Boundary regeneration preserves root provenance but is distinct from true transparent residence correction. Far-root authenticated capsules, calibrated TX/reference conversion and two-or-more-hop validation remain future work. Periodic SyncInfo and standard channel selection may narrow listening windows; Channel Sounding frequency-offset/RTT/phase observations may inspire diagnostics, not a generic passive nanosecond clock API [6, 7]. Optional UWB adapters such as DW3110 require accessible hardware and antenna-delay, NLOS, asymmetry, oscillator and energy studies [8]. Further precision is a testable matched-hardware hypothesis, not a promise.

8.3 Applications and interoperability

Home and small-business clocks, sensors and event correlation are the first intended applications. Heterogeneous passive receivers form a continuum: coarse UTC, lightweight matched follow-up, physical-stamp refinement and separately authorized diagnostics/elections. Lower-capability devices need not implement the premium estimator or transmit to consume time.

Industrial adapters may extend a selected wired PTP domain to edge timestamping and synchronization with explicit ingress profile, domain, timescale, gateway delay and uncertainty. Broadcast timecode extension requires rational frame-rate, drop-frame and epoch mapping; time labels are distinct from qualified output phase/genlock [9, 10]. Medical sensing/timestamp correlation remains exploratory after power and reliability specifications are established, with no clinical claim.

The companion protocol proposal specifies the transport registry, byte layouts, state/quality contract, exact-copy assembly, authority admission, known-answer exchanges and conformance tests. Those definitions support interoperability review. Passing codecs, demonstrating a timing mechanism and obtaining controller/SIG/IEEE qualification remain distinct outcomes.

9 Conclusion

Prospective validation shows that a separately trained 10,516 ns hardware-edge correction reduces the two clean T096 PPS-shadow populations to 70.197/76.359 ns uncentered RMS, with signed bias $-11.137/+23.992$ ns. The same fixed value works in both tested directions without validation refitting or receiver-GPS training. Their MAE, SD and P99 figures are numerically lower than MicroSync’s reported figures under the stated, different metrology.

The heterogeneous C2 result adds a service finding: all 359 normal targets are scoreable at 319.5 ns internal RMS, and serving-aware incumbent protection shortens observed fallback episodes to 2.000–3.367 s. Cold error, publication gaps and interrupted-authority tails remain. Taken together, these experiments support calibration before validation, immutable physical identity and distinct authority/publication metrics. Independent output, robust held-time behavior and measured energy are the next steps from demonstrated mechanisms to a qualified device profile.

Acknowledgments and AI-assistance disclosure

Jonathan Yantis (Engineer) provided hardware, original protocol and firmware, prior timing work, experimental direction and research decisions. OpenAI Codex assisted protocol review, implementation and test tooling, evidence analysis, literature verification, manuscript drafting and LaTeX preparation. This follows the project’s CREDITS.md; human review and responsibility for any eventual submission remain with the author. No affiliation or additional human coauthor is asserted.

Data and reproducibility

The local companion manifest binds each table/figure to exact run IDs, seals, source/image hashes, compact indexes and denominator definitions. It retains training, clean, interrupted, cold, earlier-profile and incomplete populations without pooling. Predecessor manuscript/evidence hashes remain recoverable through Git and the preserved manifests. Standalone LaTeX includes bibliography and vector graphics; build/QA manifests record source, compiler and PDF identity. Raw logs remain private. Seals prove integrity, not calibration or independent durable backup. Neither paper has been submitted or published.

References

- [1] M. Maróti, B. Kusy, G. Simon and Á. Lédeczi, “The Flooding Time Synchronization Protocol,” *SenSys*, pp. 39–49, 2004. doi:10.1145/1031495.1031501. [Author-institution paper](#).
- [2] F. Asgarian and K. Najafi, “BlueSync: Time Synchronization in Bluetooth Low Energy With Energy-Efficient Calculations,” *IEEE Internet of Things Journal*, 2022 (online 2021). [Publisher record](#). See also F. Asgarian, doctoral dissertation, University of Michigan, 2022, ch. 2, [institutional repository](#).
- [3] R. Ohara, S. Izumi, S. Imanaka, T. Yamamura, T. Ishii and H. Kawaguchi, “MicroSync: Sub-Micro Second Accuracy Wireless Time Synchronization Service,” *IEEE Access*, vol. 12, pp. 124478–124494, 2024. doi:10.1109/ACCESS.2024.3446668. [Author-university version of record](#).
- [4] S. Sridhar, P. Misra and J. Warrior, “CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth LE Advertisers,” 2015, revised 2016. [arXiv:1501.06479](#).
- [5] Nordic Semiconductor, “Wireless timer synchronization among nRF5 devices,” 2016. [Reference design and output measurements](#).
- [6] Bluetooth SIG, “The Bluetooth Low Energy primer.” [Official primer](#); Bluetooth Core Specification, Link Layer and advertising procedures. Accessed 2 October 2026.
- [7] Bluetooth SIG, Core Specification 6.3, Vol. 6, Part H, “Channel Sounding.” [Official specification](#). Accessed 2 October 2026.
- [8] Qorvo, “DW3110,” product documentation and linked data sheet; APS014, “Antenna Delay Calibration of DW1000-Based Products and Systems.” [Official product and application-note index](#). Accessed 2 October 2026. The calibration note concerns DW1000; it is methodological context, not a DW3110 calibration result.
- [9] NIST, “Meanings of common terms used in IEEE 1588.” [Clock and timing terminology](#). IEEE 1588 Working Group, [official working-group site](#).
- [10] SMPTE, ST 2059-1:2021, “Generation and Alignment of Interface Signals to the SMPTE Epoch.” [Publisher document](#). Edition and application profile must be selected for an interoperability trial.
- [11] Z. Zhou, T. Cao, C. Shen, J. Zhang, Y. Liu and H.-W. Huang, “Multiprotocol Wireless Timer Synchronization for IoT Systems,” 2026. [arXiv:2604.07199](#). Proprietary radio timeslots; stack-inactive and BLE-active conditions are distinct.
- [12] Nordic Semiconductor, hardware-team response, “Undocumented TX–RX radio delay.” [Official engineering discussion of ADDRESS decoding and TX-chain latency](#). Accessed 2 October 2026; qualitative mechanism, not this study’s calibration constant.
- [13] J. Yantis with OpenAI Codex assistance, PToB local research archive, evidence cutoff 2 October 2026. Companion `paper/evidence.json` provides run IDs, seals, source hashes, metric populations and claim mappings; `paper/README.md` provides offline reproduction instructions. Private raw captures are not distributed with these drafts.