

PToB Viability: Rooted BLE Timing Meshes for Heterogeneous Edge Networks

Protocol Architecture, Experimental Evidence, Scaling Limits,
and a Path to PTP Edge Extension

Jonathan Yantis

Independent engineering research

Research preprint for public review — 16 October 2026

Abstract

Affordable synchronization at the network edge requires more than a small clock-fit residual: it requires source provenance, causal time delivery, bounded aging, recoverable topology, and interfaces that preserve timing meaning across heterogeneous hardware. This paper evaluates Passive Precision Time over BLE (PToB) as a complete distribution architecture. PToB combines rooted source selection, authenticated grandmaster and hot-standby roles, periodic physical-event timing, regenerating relay clocks, passive clients, and a USB-to-macOS application-clock extension. We synthesize separately retained experiments rather than treating successive prototypes as one qualification campaign. A prospective two-link, one-relay experiment produced 62.504 ns fractional-second phase RMS on 436 valid causal PPS observations over 601.576 s, with 72.425% conservative usable coverage. Later evidence demonstrates bounded actual hot-standby operation, heterogeneous depth-two relay reception with unqualified transmit estimates, and a five-minute macOS run with 328.413 ns causal USB forecast RMS and 414.351 ns settled shared-source clock-agreement RMS. These measurements establish useful mechanisms, not calibrated absolute UTC accuracy or continuous service. We derive conditional hop-count, airtime, availability, holdover, and cost models; distinguish present profile limits from architectural potential; and propose a hardware-timestamped PTP ingress bridge. The strongest near-term case is a shallow, managed edge timing overlay for event correlation and sensing. Broadcast timing, deterministic industrial control, deep mobile meshes, and low-power precision operation require additional calibration, output, continuity, and energy evidence. The architecture is technically credible as a research platform and prospective economical edge extension, with explicit work remaining before application qualification.

Keywords: clock synchronization; Bluetooth Low Energy; boundary clocks; rooted mesh; hot standby; sensor networks; PTP; edge computing.

1 Introduction

Many edge devices need a shared interpretation of when an event occurred, even when they cannot justify a dedicated timing cable, a GNSS receiver and antenna at every endpoint, or a complete Ethernet timing interface. Sensors may move between coverage areas; instruments may bridge a radio clock into a general-purpose computer; relay nodes may extend coverage around machinery or through a building. These cases motivate distributing a traceable clock coordinate through an inexpensive local radio infrastructure.

The engineering question is whether a clock-distribution protocol can preserve useful time across these boundaries while remaining honest about validity and uncertainty. A device that displays a plausible timestamp after its source disappears is not necessarily delivering a timing service. Likewise, a relay that forwards an old source

timestamp without accounting for residence, local oscillator aging, and path changes can increase apparent coverage while degrading temporal meaning.

PToB addresses this question with a rooted distribution architecture. A small source-control plane selects and authorizes a serving grandmaster (GM) and, when prerequisites hold, an actually synchronized hot standby (HOT). A periodic timing plane supplies physical-event observations and subsequent corrections. Relay clocks (RCs) maintain their own causal clock maps and generate fresh downstream events. Clients consume those maps at different capability levels, including a macOS application clock reached through a USB timing bridge. The network can have redundant radio reachability while each active timing path remains rooted and acyclic.

This paper’s contribution is a system-level viability synthesis, not a claim that all pieces have already passed one end-to-end deployment test. We (i) state a transport-independent timing and authority contract; (ii) organize heterogeneous experimental evidence by what it actually measures; (iii) derive conditional scaling and energy models; and (iv) identify a practical infrastructure path from a PTP backbone to BLE edge clients. The argument deliberately transcends the prototype boards. Their measurements establish feasibility at particular boundaries; they do not define the protocol’s ultimate scope or transfer automatically to another controller.

The evidence cutoff is the completed project work available on 16 October 2026, with the latest selected physical measurements from 15 October. No new hardware measurements were performed for this synthesis. Older raw-radio, V14 fixed-route, later dynamic relay, and USB/chrony populations remain distinct. The author has approved this research preprint for public review and later submission consideration. It is not an arXiv submission or an independently reviewed certification.

2 Context and criteria for viability

2.1 Relation to existing synchronization methods

MAC-level event timestamping, skew estimation, and multihop clock dissemination have a substantial history. FTSP combines low-level timestamps and clock estimation in sensor networks [1]. CheepSync studies synchronization for resource-constrained BLE advertisers [2]. MicroSync evaluates sub-microsecond wireless synchronization with a different implementation and experimental methodology [3]. These works provide context for PToB; passive reception, regression, or later timestamp correction alone is not asserted as novel. Different reference instruments, populations, calibration procedures, error statistics, and power boundaries preclude a numerical superiority claim from unlike published results.

IEEE 1588 defines PTP for networked measurement and control, including heterogeneous clocks and application profiles [4]. SMPTE ST 2059-2 specifies a PTP profile for professional broadcast applications [5]. PToB uses analogous source-ranking and boundary-clock ideas but is a proprietary timing protocol carried over BLE, not IEEE 1588 over a new standardized medium. It is also distinct from the Bluetooth Mesh protocol. Use of BLE periodic advertising is a bearer choice, not proof of Bluetooth qualification or timing-profile conformance [6].

A rooted structure is useful independently of timing. RPL, for example, separates routing topology, rank, and repair in low-power networks [7]. PToB nevertheless needs additional clock-specific properties: an authenticated route does not establish source quality, freshness, physical event accuracy, or permissible holdover. A route repair that succeeds for data can still require the delivered clock to become unavailable during reacquisition.

2.2 A multidimensional viability test

We use five criteria. *Temporal validity* asks whether delivered time refers to an identifiable source, physical event, map generation, and permitted interval. *Precision and calibration* ask whether useful relative phase is achieved and which systematic uncertainties remain. *Availability* includes startup, expiry, reference loss, route changes,

and recovery. *Scalability* includes control state, radio occupancy, receiver scheduling, and path depth. *Economics* includes hardware, installation, energy, maintenance, and qualification cost.

A platform can be viable for timestamping intermittent sensor events while unsuitable for continuous actuator synchronization. Application requirements must therefore specify the observation point, maximum error or statistical objective, required availability, outage behavior, and trusted source. An advertised “sub-microsecond” value without these attributes is insufficient.

3 Protocol architecture and invariants

3.1 Rooted timing over a potentially redundant mesh

The logical timing graph is a tree or directed acyclic rooted subgraph selected from the available radio graph. A receiver disciplines to one selected parent in a specific source scope. Alternative parents provide recovery choices rather than clocks to average indiscriminately. The root represents independent source authority; an RC extends its time without acquiring GM entitlement. A leaf can remain receive-only during steady timing service, although admission or renewal may require occasional interaction.

Figure 1 separates demonstrated interfaces from the proposed PTP ingress. Neither a desktop application nor an unqualified relay can silently promote an uncertain local estimate to an independent time source.




Figure 1: System architecture. The relay, HOT, and host results come from separate experiments; the complete drawing is not a measured combined deployment.

3.2 Physical events and causal clock maps

A useful timestamp is attached to a named physical event, not merely a host callback. An original radio event is identified by its sender and stream incarnation, event/publication identity, and applicable channel/capture context. A later correction supplies the transmitter’s timestamp for that event. The receiver associates it with the original local capture, rejecting stale, conflicting, or ambiguously matched records.

For local counter c and immutable generation g , a generic map is

$$\hat{T}_g(c) = a_g(c - c_g) + b_g. \quad (1)$$

The implementation may use fixed-point arithmetic and a slewed actuator rather than this literal expression. The essential contract is that the selected map was already available at the time being evaluated. Fitting a line using

future observations and evaluating old points is an estimator diagnostic, not proof of a clock that was delivered then. Map handover must preserve the declared continuity policy or explicitly withdraw service.

A relay maps its actual downstream transmit event into the root coordinate and later publishes the corresponding correction. This is boundary-clock regeneration, not packet forwarding with a guessed residence delay. It permits downstream periodic service between upstream observations, but does not manufacture fresh source authority. Estimated TX events, as currently used by a heterogeneous prototype, can establish functional reception while leaving physical timing uncertainty unknown.

3.3 Four distinct lifetimes

The implementation experience strongly favors separating (1) authenticated control liveness, (2) independent source permission, (3) local timing-map validity, and (4) child service admission. A live peer can temporarily lack a precision map; a long child lease cannot extend short source permission. Freshly receiving old provenance cannot reset its age. Display smoothing cannot extend the API's clock validity.

A published sample therefore carries or binds root identity, physical boot/source epoch, selected path, sender/stream incarnation, map generation, and original deadlines. Changes fence incompatible history and cancel queued delivery where appropriate. A “HELD” output is meaningful only within a declared valid hold interval. Missing, unknown, withdrawn, and expired states remain distinct; none becomes valid because the previous RMS was small.

For example, the fixed-route research profile uses local map age strictly below 500 ms. Its admission representation can last up to 360 s, while source permission is much shorter [10]. These are different clocks in the protocol state machine, not alternative precision guarantees. Inherited source age must continue to accumulate across regenerating relays.

3.4 Source selection and hot standby

Source eligibility precedes ranking. A node needs independent source evidence and appropriate admission before becoming a candidate; good mesh-derived time cannot make it independent. PTP-inspired ordering does not eliminate the need to handle asymmetric visibility, new arrivals, and partitions. A current implementation principle is that a newcomer yields without destabilizing a demonstrably live established GM/HOT pair. Loss of the precision map alone is not evidence that the primary has died.

HOT is an actual operating condition: the selected standby is synchronized and has current takeover permission. A standby label, capable hardware, or an election preference is insufficient. The purpose is to reduce recovery work by preparing a replacement, not to promise zero phase discontinuity or immediate failover. Independent source loss, takeover fencing, source incarnation, and downstream reacquisition still matter.

Disconnected partitions cannot infer a globally unique GM from local radio evidence. They may operate distinct source scopes, then reconcile on restored visibility. The present evidence does not establish arbitrary partition/merge convergence or three-GM stability. Production deployment should bound the source-control domain and avoid turning every client or relay into a GM participant.

3.5 Authentication and heterogeneous capability

The frozen V14 design uses truncated HMAC-SHA256 with provisioned group credentials and separate authentication domains for origin and local records [10]. This provides group authentication, not a publicly verifiable signature or protection against a compromised member able to forge the same domain. Boot/epoch, sequence, admission challenge, path, and freshness checks address replay and stale interpretation. Cryptographic authenticity does not prove physical time accuracy and does not prevent RF jamming.

Heterogeneity should be explicit in capabilities: actual versus estimated event capture, counter resolution, clock-map availability, calibrated delay, output interfaces, oscillator control, and uncertainty state. A coarse display client, an actual-event receiver, a software-estimated transmitter, and a hardware-timed output device can share source provenance without sharing a precision claim. Radio adapters must preserve supported controller ownership, event identity, and bounded work; moving a timestamp into application code is not automatically equivalent to a hardware capture.

4 Evidence synthesis and measurement boundaries

4.1 Method

We reviewed existing manuscripts, result records, protocol summaries, and their indexed provenance. The companion evidence manifest hashes the selected inputs and maps claims to them. This paper does not reanalyze all private raw archives or certify every earlier experiment. No unlike runs are pooled. We distinguish uncentered RMS (including signed bias), demeaned jitter, causal forecasts, estimator-reported RMS, and independent output measurements. Status observations are not continuous-time coverage.

Table 1 summarizes the strongest complementary results. The central conclusion is compositional: several necessary interfaces function, while their joint behavior and uncertainty budget remain to be validated.

Table 1: Separate evidence populations. Each limitation is part of the result.

Boundary	Observation	Evidentiary limit
One relay, two BLE links	601.576 s; 436 valid causal PPS samples; 62.504 ns RMS; 191 ns maximum	436/602 conservative usable boundaries; fractional-second phase; absolute UTC and output-pad uncertainty unknown
GM / actual HOT	24/24 qualified paired status rows over 178.621 s	Bounded paired capability; physical loss-to-output takeover not instrumented; three-GM stability unresolved
Heterogeneous relay	Named RC2 paths followed in 91/116 frames for each of two separate treatments	Functional authenticated depth-two reception; estimated TX; precision unqualified
USB causal mapping	9437 pre-ingest inverse forecasts; 328.413 ns RMS	Prediction of observed USB tuples; unknown host/PHY bias and no independent UTC reference
macOS delivered clock	Settled 60–300 s: 414.351 ns shared-source agreement RMS; 182 valid / 54 withheld	Shared inputs; 1315 ns maximum; not independent UTC accuracy or continuous availability

4.2 Native clock discipline and one-relay timing

Earlier software-oscillator experiments established causal running-clock behavior, while also demonstrating sensitivity to operating conditions. Separate quiet-pair captures measured 69.136 and 68.106 ns causal RMS on 1144/1144 and 727/728 retained originals. A simultaneous three-unit capture measured 1075.679 ns on 354/384 scoreable originals and never reached qualified native lock among 59 status observations [9]. These raw-radio-era data are historical mechanism evidence, not interchangeable measurements of the later periodic controller profile.

The stronger relay experiment used a source GM, one dedicated RC, and a leaf evaluated at its independent local PPS witness. The leaf’s witness was excluded from its delivered estimator and peer acquisition aid. A direct calibration on a separate population fixed a provisional +10543 ns receive correction per BLE link. It was

applied once per received link before the prospective relay run; the two-link correction was therefore +21086 ns [11, 10].

Over 601.5761865 s, 504 physical PPS records were retrieved. There were 436 reference-valid causal comparisons, with mean -18.417431 ns, uncentered RMS 62.503871 ns, and demeaned jitter 59.728822 ns. Absolute P95/P99/maximum were 123.25/160.25/191 ns. Sixty-five observed edges lacked valid witness metadata and three reference-valid edges had an expired delivered map. Using 602 conservative nominal boundaries yields 72.425% usable coverage, including unobserved boundaries. The largest witness gap was 59.017190 s. Separately, sampled precision accounting recorded six relay and eleven leaf map expiry/recovery episodes; sparse PPS intersections cannot erase those gaps.

This is a compelling fractional-second phase result for a causal two-link service with calibration frozen beforehand. It does not identify the physical cause of the empirical correction, calibrate receiver/cable delays, establish the UTC second association, or measure a delivered electrical output. Calibration portability across directions, boards, channels, temperature, and geometry is unknown. The native 62.5 ns counter grid does not prevent statistical sub-tick map coefficients, but neither does such estimation create a finer hardware output edge.

4.3 Root and HOT capability, including regressions

The selected 15 October GM checkpoint observed actual GM/HOT pairing in all 24 sampled rows over 178.621 s, with no sampled competing serving claims. A later separate monitor recorded a serving primary in 80/80 rows and HOT recovery in 64/76 standby rows [12]. These data establish that actual synchronized standby is achievable, rather than only represented in a model or UI.

The same checkpoint explicitly retains unresolved instability with three concurrent GM-capable clocks. The accepted working arrangement reduced the third clock to RC operation. Operator-reported reboot success was not accompanied by an instrumented GM restart and takeover in the retained monitor, so it is not a failover latency or phase-continuity measurement. Later GNSS validity loss caused source withdrawal and readmission; a separate 179.717 s capture retained 19 no-serving rows among 60 and no actual HOT [13]. The source-validity policy matters as much as the radio fit: redundancy requires a qualified surviving source and correct control behavior.

The viability claim is consequently bounded. A small rooted domain can support serving and actual HOT roles. General multi-GM stability, controlled failure recovery, uninterrupted output, and merge behavior remain research obligations. Retained regressions reveal the limits of inferring reliability from a favorable steady-state snapshot.

4.4 Heterogeneous relay and dynamic service

The 14 October demand-driven W12 release exercised actual reception through native and heterogeneous RCs [14]. Separate named-parent treatments observed direct RC1 following in 108/117 and 106/115 frames. In two separate RC2 treatments, each receiver followed the expected authenticated depth-two path in 91/116 frames. Old-parent, rebind, SEARCH, HOLD, and unavailable observations remain in those denominators.

These are functional relay results. Transmit tick/channel values remain estimates with unknown bounds, and no T096-through-W12 PPS precision result was established. It would be invalid to assign the native 62.504 ns distribution to these paths. The experiments nevertheless demonstrate that origin/path semantics, request/response admission, downstream timing reception, and heterogeneous forwarding roles can operate together.

The timing train is demand-driven: discovery remains available, but timing publication requires live authenticated child demand and a usable upstream map. One natural child expiry was consistent with the original 360 s lease; renewal uses the existing 300 s interval. The idle relay retained discovery and upstream reception. This avoids unnecessary timing publication structurally, but controller state is not an RF silence measurement and no energy saving was measured. Observed source changes also produced HOLD and recovery without W12 reboot. They do not establish arbitrary mobility or long-term availability.

4.5 Extension into macOS

The USB bridge associates device counter captures with host frame/Mach observations using causal markers and immutable source maps. This avoids treating USB request arrival time as the physical timing event. The fine application-clock path composes the USB map with the retained source coordinate and feeds a private `chronyd -x`; the operating-system clock is not steered. The `chrony` documentation describes the `no-system-clock-control` option [8]; the project result concerns the private application path actually exercised.

The final 300.065 s run used TC04 firmware 0.15.53 and CLI 0.21 [15]. Full-run pre-ingest inverse forecasts measured 328.413 ns RMS across 9437 qualified tuples. The distinct forward innovation population measured 327.448 ns RMS across 9494 observations. During the fixed settled 60–300 s window, 265 new `chrony` queries reported estimator RMS of 23–339 ns, median 129 ns. Separate delivered-clock comparison against the same-instant fine source estimate measured 414.351 ns RMS on 182 valid observations, with 54 withheld and a 1315 ns maximum. These are three different metrics, not multiple independent estimates of absolute accuracy.

The full startup transient reached 32.072201 ms. The run recorded zero USB errors and zero actual private-daemon restarts, compared with 46 restarts in a separate previous 300 s run. No natural execution of the newly distinguished missing-SOF-map retention branch, refusal, or original-expiry condition occurred; synthetic branch coverage is separate. The sequential improvement therefore does not prove that change caused the lower restart count. Finite helper completion at 299.990533 s and ensuing cleanup remain distinct from an unplanned source outage.

The result shows that useful relative timing can survive an inexpensive receiver-to-computer interface. It does not calibrate initial UTC offset or host/PHY bias, qualify the macOS system clock, or promise deterministic application scheduling. A timestamp query and a physical output action on a general-purpose OS have different latency and uncertainty budgets.

5 Infrastructure topologies and extensibility

5.1 Managed cells and shallow relay trees

A practical initial deployment is a managed timing cell: one independent root, one qualified standby where justified, a small set of powered RCs, and many passive clients. A star minimizes accumulated link uncertainty. A shallow relay tree extends around coverage obstacles while retaining a simple selected parent and diagnosable provenance. Powered infrastructure and battery endpoints can have different duty policies.

A long corridor can use a chain, but depth increases age, bias exposure, correlated failure, and scheduling constraints. Where Ethernet or another qualified source distribution already exists, several shallow BLE cells may be preferable to one deep chain. Common upstream provenance allows their time relationship to be analyzed; it does not by itself prove phase alignment across bridges.

Multiple candidate parents can support mobility and local repair, provided switching fences the old path, reacquires the new one, and preserves the original source permission. RSSI and stream availability can guide attachment but are not timing-quality measurements. Parent dwell and hysteresis can reduce churn; a newly selected path may require different calibration. Fast packet connectivity must not masquerade as continuous clock service.

5.2 Profile limits are separate from architectural reach

The frozen V14 fixed-route contract contains at most eight route entries, three distinct RC positions, and at most two consecutive RC intermediaries. Its GM-control domain is capped at four directly connected participants including the coordinator [10]. Those are bounded implementation/specification limits, not evidence for four

stable GMs or an eight-node precision deployment. Later RC7/dynamic selection work is a successor profile and does not retroactively make the fixed-route contract dynamic.

For clarity, we define H as the number of received BLE links from root to endpoint. A direct leaf has $H = 1$, one RC plus a leaf has $H = 2$, and two RC intermediaries plus a leaf has $H = 3$. The precision experiment validates the selected $H = 2$ path. Heterogeneous depth-two reception is functional evidence only. Analytical values for larger H below describe a possible future generalized profile, not permission to exceed current route or controller limits.

5.3 A transport-neutral clock interface

Extensibility is most useful at the boundary between a clock contract and a transport adapter. The contract should return time, source scope, local and inherited age, uncertainty status, validity deadline, and discontinuity information. An adapter supplies actual physical-event evidence and its capture semantics. USB, Ethernet ingress, alternate MCUs, and future output hardware can then be evaluated against the same obligations without claiming identical performance.

An optional tunable oscillator or synthesizer could move from software clock prediction to physical frequency/phase control. That requires measured actuator gain, delay, quantization, saturation, and a safe control loop. Current software-map convergence does not demonstrate quartz discipline, carrier coherence, or calibrated PPS output. Similarly, LoRa, WiFi coexistence, and long quiet windows need their own timing and resource validation; they are not established by the BLE results summarized here.

6 Conditional hop-count and scalability analysis

6.1 Random error, systematic bias, and common source error

Let the endpoint error after H links be

$$e_H = e_0 + e_{\text{out}} + \sum_{i=1}^H (b_i + \varepsilon_i) + e_{\text{age}}, \quad (2)$$

where e_0 is source error, e_{out} is the final interface/output error, b_i is residual link bias, ε_i is zero-mean link variation, and e_{age} is holdover error. With total mean μ_H , the general RMS includes covariance:

$$\text{RMS}_H^2 = \mu_H^2 + \text{Var} \left(e_0 + e_{\text{out}} + e_{\text{age}} + \sum_i \varepsilon_i \right). \quad (3)$$

Only under uncorrelated terms does the variance reduce to a sum. Shared source errors, calibration offsets, temperature, channel effects, and estimation histories can be correlated. Differential timing between two leaves may cancel common source error, but it does not cancel their distinct branches or output delays.

To anchor sensitivity to the measured scale without pretending to identify per-link noise, define the *hypothetical effective increment*

$$s_{\text{eff}} = 62.503871 / \sqrt{2} = 44.1969 \text{ ns}. \quad (4)$$

This allocation treats the entire observed two-link RMS as two independent equal terms, with zero source, output, bias, and age contributions. Those assumptions are unverified. It is neither a measured per-link parameter nor a conservative confidence bound. Two illustrative extensions are $E_{\text{ind}}(H) = s_{\text{eff}}\sqrt{H}$ and $E_{\text{corr}}(H) = s_{\text{eff}}H$, the latter representing perfectly aligned equal random contributions at that assumed scale. It is a sensitivity alternative, not a second fit to the two-link result.

Table 2: Hypothetical RMS scaling in ns. Only the $H = 2$ independent allocation is anchored to the measured RMS by construction; other entries are projections.

BLE links H	RC intermediaries	Independent model	Fully correlated model
1	0	44.2	44.2
2	1	62.5	88.4
3	2	76.6	132.6
4	3	88.4	176.8
8	7	125.0	353.6
16	15	176.8	707.2

For an illustrative RMS allocation $B = 500$ ns, ignoring every other term gives

$$H_{\text{ind}} \leq \lfloor (B/s_{\text{eff}})^2 \rfloor = 127, \quad H_{\text{corr}} \leq \lfloor B/s_{\text{eff}} \rfloor = 11. \quad (5)$$

These deliberately contrasting arithmetic ceilings are *not supported hop counts*. RMS is not a worst-case guarantee, 500 ns here is an illustrative budget, and neither result qualifies the separate project requirement of bounded error throughout a blackout. If a hypothetical same-sign residual bias of 100 ns per link is added, the independent model becomes $\sqrt{Hs_{\text{eff}}^2 + (100H)^2}$ ns: four links give about 409.7 ns while five give 509.7 ns. Small persistent bias can therefore dominate before random scatter does. No 100 ns residual bias has been measured as a universal link parameter.

Unknown calibration and output terms must remain unknown, rather than being assigned zero in a deployment budget. The two-link experiment cannot uniquely separate source, reference, capture, and relay components. Deeper independent experiments are needed before choosing even the independent versus correlated model, much less a maximum depth.

6.2 Source age and scheduling

A relay's local update resets only its local estimator age. A simple age accounting model is

$$A_H = A_{\text{ingress}} + \sum_{i=1}^H d_i, \quad (6)$$

with each d_i including the relevant capture-to-correction availability, mapping, and forwarding wait. These values require measurement. At a 16 Hz timing cadence the period is 62.5 ms. If, purely illustratively, each link consumed a full period, eight links would consume 500 ms before ingress age, missed events, or margins. This is not the measured per-hop delay, and the present local 500 ms rule is not an end-to-end age guarantee. Delayed corrections can span more than one event; off-phase scheduling can reduce some waits.

Slot capacity must satisfy a collision-domain constraint such as

$$K(t_{\text{event}} + t_{\text{guard}}) + t_{\text{control}} + t_{\text{reserve}} \leq T_f. \quad (7)$$

Here K is active transmitting trains in the domain, not passive receivers. Controller activity limits, up-stream/downstream receive conflicts, coexistence, and guards may bind before packet duration does. Spatial reuse requires evidence about interference, hidden terminals, and channel scheduling; multiplying a single-cell capacity by area is unjustified without it.

6.3 Airtime, fan-out, and control load

A periodic broadcast can serve many listeners without one precision packet exchange per listener. For the 190-byte research record at 16 Hz, the record bits alone require $190 \times 8 \times 16 = 24320$ bit/s, or 2.432% of a 1 Mbit/s raw

link rate. This is a lower-bound serialization calculation, not measured RF duty: it excludes PHY/link overhead, discovery, auxiliary scheduling, retransmission/redundancy, control, and guard time. Ten fully overlapping trains would consume at least 24.32% on that simplistic serialization basis, without demonstrating a feasible schedule or collision performance. The 212-byte successor format and other profiles require their own accounting.

Steady broadcast TX work is approximately independent of passive client count, but admission state and renewal work are not. With N_a admitted children and renewal interval T_r , a first-order average request rate is N_a/T_r . For 1000 children and 300 s renewals this is 3.33 renewals/s before retries or bursts. This hypothetical number is not within demonstrated controller or table capacity. Admission serialization, memory bounds, simultaneous startup, and source changes can dominate even when precision airtime is modest. Bounded control domains and hierarchical cells are more credible than an unlimited shared admission server.

6.4 Availability scales differently from RMS

If each link were independently usable with probability p at a defined observation instant and every link were required simultaneously, path usability would be p^H . For $p = 0.99$, ten and twenty links give 90.4% and 81.8%; for $p = 0.999$, twenty give 98.0%. These are illustrative probabilities, not estimates from the selected experiments. Holdover changes the event definition and correlated outages invalidate independence. The measured 72.425% relay scoreability includes reference validity and cannot be substituted for a per-link packet-success probability.

Redundant parent choices may improve coverage but incur acquisition and continuity costs; two correlated paths are not independent redundancy. Source power loss, interference, a shared GNSS obstruction, or a common software defect can affect all routes. A viable network must report both the conditional error distribution when usable and the time/observation denominator over which it is usable. The effective deployment depth is limited jointly by timing, age, schedule, availability, trust, and supported profile capacity.

7 Power sampling, accuracy, and economics

7.1 Why the periodic stream is well positioned

Periodic timing offers predictable reception opportunities, repeated frequency observations, and a way to amortize acquisition and clock fitting. Multiple clients can share a stream, and a receiver can potentially accumulate sufficient observations in a scheduled window before sleeping. Complete corrections avoid relying on incomplete fragment assembly for each useful update. Demand-driven relay publication can avoid serving an empty branch.

These properties make the architecture promising for energy-aware synchronization, but the experiments do not yet establish the optimum sampling rate or energy per valid timestamp. The current heterogeneous relays in the selected field release remain BLE100. A half-duty client objective, an idle timing advertiser, and a low-power product claim are different things. CPU scheduling, display, GNSS, regulator loss, scanning acquisition, retries, and logging can dominate whole-device energy.

7.2 An explicit holdover and energy model

For a proposed sleep interval T , a deterministic budget could be

$$E(T) \leq E_{\text{start}} + B_y T + \frac{1}{2} B_{\dot{y}} T^2 + E_{\text{readout}}, \quad (8)$$

where B_y bounds residual fractional-frequency error and $B_{\dot{y}}$ bounds its change. These must be experimentally bounded, including temperature and aging. With a hypothetical remaining allocation of 500 ns and no drift term, 1 ppm permits 0.5 s, 0.1 ppm permits 5 s, and 10 ppb permits 50 s. None is an observed safe sleep interval or authority extension. Source permission may expire before an oscillator budget is exhausted.

A first-order power model is

$$\bar{P} = DP_{\text{rx}} + (1 - D)P_{\text{sleep}} + f_{\text{wake}}E_{\text{acq}} + P_{\text{compute}} + P_{\text{other}}. \quad (9)$$

At short windows, repeated synchronization acquisition can erase radio-off savings. Fewer samples can reduce estimator quality; excessive filtering can hide temperature-induced rate changes. Conversely, always listening need not improve accuracy if dominant bias is unchanged. The relevant optimization is energy per application-usable event subject to error, availability, and source-age constraints, rather than duty cycle alone.

The next study should sweep reception window, update rate, source/relay role, temperature, RF load, and motion while independently measuring current and delivered error throughout the entire blackout. It must retain failed acquisitions, wake latency, all unavailable intervals, and observer effects. A prospective held-output trace is required; retrospective fitting across future samples cannot qualify a quiet interval. This is especially important for the project's deferred ± 500 ns total held-time objective at at most 50% listening.

7.3 Economics without unsupported prices

The economic hypothesis is that a small number of infrastructure clocks can distribute useful time to existing BLE-capable endpoints, reducing endpoint GNSS hardware, antennas, cabling, or specialized network interfaces. It is not yet a measured cost advantage. A complete model is

$$C_{\text{total}} = C_{\text{sources}} + C_{\text{relays}} + NC_{\text{endpoint}} + C_{\text{install}} + C_{\text{energy}} + C_{\text{maint}} + C_{\text{qual}}. \quad (10)$$

BLE silicon availability alone does not make precision integration free: controller access, stable oscillators, metrology, calibration, secure provisioning, and maintenance have costs. A two-level product strategy is plausible: low-cost coarse or event-correlation clients and selectively calibrated precision endpoints. The decision should compare cost per qualified usable endpoint against application-specific wired or network-timing alternatives, rather than quote an unmeasured universal savings percentage.

8 A future PTP-to-PToB edge bridge

The most direct extension to existing infrastructure is a PTP-disciplined ingress clock that publishes PToB time into a local BLE cell. A hardware-timestamped Ethernet interface would discipline a local clock, and a calibrated mapping would connect that clock to the radio capture/transmit counter. The PToB root would then advertise the upstream source identity and validity conservatively. This is a proposed gateway architecture; no PTP bridge is demonstrated in the selected evidence.

The bridge must preserve the time domain, timescale, epoch, UTC offset and leap handling, traceability/quality state, source changes, and loss behavior. An upstream GM change must not silently reuse incompatible phase calibration or stale source authority. The Ethernet-to-radio cross-timestamp must be measured with its latency and uncertainty; reading an OS wall clock after a packet callback is inadequate for a precision ingress claim. Independent source policy must also recognize that two bridges sharing one upstream PTP source do not provide independent source redundancy.

For a deterministic end-to-end bound, a proposed allocation is

$$B_{\text{edge}} \leq B_{\text{PTP}} + B_{\text{cross}} + \sum_i B_{\text{BLE},i} + B_{\text{hold}} + B_{\text{output}}. \quad (11)$$

Every component requires a supported bound. Root-sum-square combination is appropriate only for explicitly modeled random components and does not remove bias. If any required bound is unknown, the total remains unknown. A calibrated PTP backbone cannot by itself calibrate the wireless edge.

The first implementation should be a one-way timing gateway with explicit source demotion on upstream loss, not a claim that BLE packets are IEEE transparent-clock messages. Reverse export into a PTP network would require a separate conformance and loop-prevention design. A mesh-derived clock must not be reintroduced upstream as an independent superior source. Selection of IEEE edition, amendments, profile, delay mechanism, and time properties belongs to a specific future interoperability project.

Validation should proceed from calibrated PTP-to-radio-counter mapping, to direct BLE delivery, to one relay, to source switching and held operation. Simultaneous independent measurements at ingress and edge output should establish both common-mode and differential error. A future SMPTE deployment must separately validate the media profile and physical media-clock/output requirements; this paper establishes no ST 2059 or genlock conformance.

9 Application viability

Table 3: Application-oriented assessment. Thresholds must come from the intended application, not from a generic industry label.

Application	Plausible value	Remaining decisive evidence
Industrial monitoring	Correlating machine, vibration, environmental, and maintenance events without timing cabling at every sensor	Sensor capture-to-time calibration; RF coexistence; thermal drift; specified outage/availability and service life
Industrial control	Potential shared schedule or physical output timing	Deterministic worst-case error and latency, calibrated actuation, safe failure response, continuous qualification
Broadcast support	Timestamped telemetry, production metadata, distributed event logs and auxiliary sensors around a PTP facility	Bridge calibration, source/time-property mapping, independent endpoint accuracy and operational recovery
Broadcast media timing	Prospective media-clock or output synchronization extension	Applicable PTP/media profile, genlock/sample-clock output, phase continuity and redundancy validation
Emerging sensor networks	Dense receive-oriented time sharing, mobile instruments, distributed measurement and edge gateways	Power/error Pareto measurements, mobility and parent switching, secure deployment, scalability under churn

Event correlation is the strongest initial use case because an application can explicitly reject unavailable or uncertain timestamps and retain source provenance. Higher-integrity systems should propagate quality with the event data, rather than converting all samples to an apparently uniform UTC field. Two events with different roots or source epochs require a qualified relationship before their timestamps can be compared precisely.

The broadcast opportunity is complementary to a facility timing backbone: extending time to otherwise disconnected accessories and sensors. A relative software clock suitable for log correlation is not a media sample clock. Likewise, the industrial opportunity spans very different requirements; success for condition monitoring cannot be promoted to motion-control readiness. Sensor localization or distributed acoustic processing may require differential phase and hardware capture beyond simple timestamp agreement.

A heterogeneous mesh is economically attractive precisely because all clients need not implement the most expensive timing tier. That benefit depends on honest capability discovery and quality-preserving interfaces. A coarse client must not advertise itself as calibrated merely because it recognizes the same root as a precision client.

10 Limitations and a validation program

The evidence comes from a small set of laboratory and field-oriented prototypes, short finite treatments, different software generations, and incompletely characterized geometry and environment. Some results are build-owned functional checks; the full independent suite remains held. There is no large-scale deployment, lifetime energy study, controlled mobility campaign, adversarial security validation, or calibrated complete source-to-electrical-output budget. One-way empirical compensation leaves propagation and detector/reference effects entangled. At roughly 3.3 ns per metre of free-space path, changing geometry can matter at a tens-of-nanoseconds objective; this physical scale is not a measurement of the tested path.

The present synthesis also does not prove that the best native relay, HOT, heterogeneous, and macOS results coexist in one configuration. Firmware regressions, receiver source loss, console gaps, and incomplete captures remain in the source records. Hashes bind evidence identity, not measurement truth or off-host durability. The reproducibility package accompanying this paper contains selected claim inputs and model derivations, not unrestricted private raw logs.

A staged validation program follows directly from the viability model:

1. **Metrology:** freeze calibration before prospective tests; measure named physical output edges and sensor capture points against an independent calibrated instrument; characterize source, cable, detector, and counter uncertainty.
2. **Composed operation:** exercise a qualified root/HOT pair, a native relay, a heterogeneous relay, and host delivery in a single declared topology. Retain startup, source loss, old-map rejection, and all outages.
3. **Dynamic control:** inject controlled primary loss, asymmetric reception, new arrivals, restart, partition, and merge. Score temporary dual serving, valid HOT tenure, convergence time, output discontinuity, and lost service separately.
4. **Depth and density:** progress from direct to one and two intermediaries before proposing deeper profiles. Vary active transmitters independently of passive listener count; measure RF occupancy, controller limits, acquisition, and admission bursts.
5. **Energy and environment:** pair current measurements with continuous held-time errors through scheduled quiet windows, temperature, coexistence, and mobility. Evaluate precision per joule and unavailable-event cost.
6. **PTP ingress and application qualification:** validate the gateway cross-clock map, time properties and failover; then test the named industrial or broadcast output requirement using its applicable profile and instrumentation.

These are proposed experiments, not completed work. They favor a sequence of falsifiable boundaries over a single favorable RMS target. A protocol remains useful when it refuses unsafe output; availability must nevertheless meet the application's needs rather than being dismissed as an implementation detail.

11 Conclusion

PToB has progressed beyond an isolated clock-fit experiment. The retained evidence demonstrates a causal two-link timing path with small conditional fractional-second phase error, bounded actual HOT operation, functional heterogeneous relay reception, demand-driven service behavior, and relative time delivery into macOS. Together these support the viability of a rooted BLE timing architecture and identify a credible role at the edge of existing timing infrastructure.

The evidence supports shallow managed timing cells more strongly than deep autonomous precision meshes. Broadcast fan-out is favorable, but relay depth couples bias, source age, availability, and controller scheduling. Periodic precision streams offer a sound basis for future sampling/power optimization, while current data do not yet establish energy savings or a safe low-duty precision profile. A future hardware-timestamped PTP ingress can reuse the provenance and clock-map contract, provided its own calibration and source semantics are validated.

The research position is therefore specific: an affordable, heterogeneous time-sharing overlay is technically credible and experimentally grounded at several critical interfaces. Its transition to industrial or broadcast qualification depends on integrated continuity, independent output accuracy, dynamic topology, and energy evidence. Reporting these boundaries alongside the achieved precision makes the viability case both useful and testable.

Author contributions and AI-assistance disclosure

Jonathan Yantis supplied the original protocol and firmware, hardware, experimental direction, engineering decisions, and research objectives. OpenAI Codex assisted the underlying project’s implementation, tooling and evidence analysis, and assisted this manuscript’s synthesis, literature checking, analytical models, LaTeX preparation, and document verification. No AI system is named as a human author. The human author retains responsibility for reviewing all claims and any future submission. This preprint invents no institutional affiliation, funding award, submission identifier, or peer-review status.

Data and source availability

The standalone LaTeX embeds its bibliography and architecture figure. A companion local review package contains the manuscript, derived analytical CSV tables, selected-input SHA-256 manifest, and derivation script. Public project material is available at <https://ptob.dev/>; cited internal evidence is identified below by stable repository paths and date, not represented as a deposited public dataset. Private provisioning information and raw operational captures are excluded. Durable public data deposition, licensing, and venue-specific submission checks remain to be completed before external submission. This paper does not assign a new license to inherited software or data.

A Evidence map and reproducibility

Table 4: Primary local anchors. The companion manifest records full hashes; these records retain links to underlying captures and review indexes.

Claim family	Repository anchor
Clock discipline	<code>paper/oscillator-p11-evidence-20261003.json</code>
Fixed-route architecture	<code>paper/ptob-v14-evidence.json</code> ; <code>paper/ptob-v14-preprint.tex</code>
One-relay measurement	<code>docs/RC_ONEHOP_PPS_RESULT_20261012.md</code>
HOT and regressions	<code>results/T096_GM_HOT_FIELD_20261015.json</code> ; <code>results/T096_GM.UTC_FIELD_20261015.json</code>
Heterogeneous relay macOS / USB	<code>docs/W12_RELAY_DEMAND_RSSI_RELEASE_20261014.md</code> ; corresponding result JSON <code>results/MACOS_SOF_MAP_FIVE_MINUTE_20261015.json</code> ; <code>paper/usb-macos-evidence-20261015.json</code>
Analytical models	<code>paper/figures/viability-20261016/derive.py</code> ; generated CSVs

All analytical hop, availability, serialization, and holdover values are labeled projections and can be regenerated without hardware. The effective-increment allocation uses the unrounded 62.503871 ns relay RMS. The model never estimates an unknown bias from that number. Literature was checked against the primary sources listed below on 16 October 2026. Standards catalog and publisher descriptions support context only; this paper makes no clause-by-clause conformance finding.

References

- [1] M. Maróti, B. Kusy, G. Simon, and Á. Lédeczi, “The Flooding Time Synchronization Protocol,” *SenSys*, 2004, pp. 39–49. doi:10.1145/1031495.1031501. https://archive.isis.vanderbilt.edu/sites/default/files/Maroti_M_11_3_2004_The_Floodi.pdf.
- [2] S. Sridhar, P. Misra, G. S. Gill, and J. Warrior, “CheepSync: A Time Synchronization Service for Resource Constrained Bluetooth Low Energy Advertisers,” *IEEE Communications Magazine*, vol. 54, no. 1, pp. 136–143, 2016. <https://arxiv.org/abs/1501.06479>.
- [3] R. Ohara et al., “MicroSync: Sub-Micro Second Accuracy Wireless Time Synchronization Service,” *IEEE Access*, vol. 12, pp. 124478–124494, 2024. doi:10.1109/ACCESS.2024.3446668. <https://da.lib.kobe-u.ac.jp/da/kernel/0100491658/>.
- [4] IEEE, *IEEE 1588-2019: IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, 2020. Official catalog and profile scope: <https://standards.ieee.org/ieee/1588/6825/>.
- [5] SMPTE, *ST 2059-2:2021: SMPTE Profile for Use of IEEE-1588 Precision Time Protocol in Professional Broadcast Applications*. Publisher repository: <https://github.com/SMPTE/st2059-2>.
- [6] Bluetooth SIG, *Bluetooth Core Specification 6.1*, Vol. 6, Part B, Link Layer Specification, fixed-edition reference. <https://www.bluetooth.com/wp-content/uploads/Files/Specification/HTML/Core-61/out/en/low-energy-controller/link-layer-specification.html>.
- [7] T. Winter et al., “RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks,” RFC 6550, 2012. <https://www.rfc-editor.org/rfc/rfc6550.html>.
- [8] chrony project, *chronyd(8)*, version 4.6 documentation, option -x. <https://chrony-project.org/doc/4.6/chronyd.html>.
- [9] J. Yantis, *Causal Software Clock Discipline over Wireless Physical Events: Native Lock Evidence, Operating Tiers, and a Tunable-Oscillator Extension*, local research manuscript and evidence manifest, 3 October 2026. paper/oscillator-pll-research-20261003.tex.
- [10] J. Yantis, *Passive Precision Time over BLE: V14 Boundary Clocks, Causal Timestamping, and Prospective One-Relay Validation*, local preprint and evidence manifest, 12 October 2026. paper/ptob-v14-preprint.tex.
- [11] PToB project, *TC01 source to TC02 relay to TC03 shadow PPS*, retained result, 12 October 2026. docs/RC_ONEHOP_PPS_RESULT_20261012.md.
- [12] PToB project, *GM/HOT field checkpoint*, 15 October 2026. results/T096_GM_HOT_FIELD_20261015.json. Accepted checkpoint with unresolved three-GM limitation.
- [13] PToB project, *GM UTC field investigation*, 15 October 2026. results/T096_GM.UTC_FIELD_20261015.json.
- [14] PToB project, *W12 demand-driven relay field release*, 14 October 2026. docs/W12_RELAY_DEMAND_RSSI_RELEASE_20261014.md; associated result JSON. Functional evidence; TX precision unqualified.
- [15] PToB project, *SOF-map five-minute USB/private-chrony run*, 15 October 2026. results/MACOS_SOF_MAP_FIVE_MINUTE_20261015.json. See also J. Yantis, *USB/macOS time-transfer preprint*, public chrony revision: <https://ptob.dev/papers/usb-macos-time-transfer-chrony-20261015.pdf>.