

Beyond USB Packet Arrival: Frame-Identified Time Transfer to macOS

Hardware SOF Capture, Causal Clock Mapping,
and the Separation of Stability from UTC Accuracy

Jonathan Yantis (Engineer)

Research preprint for local review — 15 October 2026

Abstract

USB is commonly used to transport timestamps, yet application-visible packet timing mixes device time with host scheduling and transfer delay. We investigate a different role for full-speed USB: a shared physical-event reference for mapping an externally disciplined device counter into the macOS Mach clock domain. A Nordic nRF52840 captures start-of-frame (SOF) events through a hardware interconnect into a free-running 16 MHz timer. A userspace IOUSBHost client reads Apple’s microframe/time reference, while nonce-bound control transactions constrain the integer relationship between host microframes and device frames independently of phase fitting. In one 600.013 s direct-XHCI experiment, 19,093 boundary tuples and 18,332 source-monotonic observations passed their respective checks. Across all twenty disjoint 30 s windows, median and worst-window detrended phase standard deviations were 0.155 and 0.288 μ s. A stricter offline test, fitting each 30 s block and predicting the following block without refitting, produced 1.053 μ s RMS and a 3.662 μ s maximum relative phase error. A single ten-minute fit instead left 3.956 μ s standard deviation with +4.635 ppm drift. Reported UTC phase retained 1.465 ms variation, dominated by the source’s UTC translation. We describe frame identity, causal map publication, permission expiry, unavailable observations and reproducible metric definitions, and compare the method with macOS gPTP, calibrated PTP and prior USB synchronization. The evidence supports short-term relative time transfer on one setup; absolute UTC accuracy, calibrated host timestamp bias and superiority to NTP remain unestablished.

1 Introduction and contributions

A timestamp can arrive late without being wrong about the event it describes. Conversely, a timely USB response can carry a poor clock coordinate. This distinction matters when bringing an embedded reference clock to a desktop application: measuring the response callback mostly characterizes the delivery path, whereas measuring a physical event in both clock domains can characterize the clocks themselves.

This study connects a reference-clock device, TC-04, to an Apple silicon Mac over full-speed USB. TC-04 is disciplined by a direct upstream PToB source, TC-01 [1]. The upstream link remains BLE; the *Mac client* uses USB exclusively. The device counter, its upstream source-monotonic map, the source UTC translation, Apple’s Mach clock and the Mac’s wall clock are therefore distinct coordinates. The experiment measures their relationships rather than assuming that one successful transport establishes all of them.

Full-speed USB supplies nominal 1 ms frame boundaries. That cadence is not an accuracy floor when each boundary is captured precisely. The difficult questions are which device event corresponds to which host event, how their timestamps relate to physical wire time, and whether the source’s UTC association is valid. A mistaken whole-frame association can shift time by 1 ms while leaving an excellent affine fit. A variable UTC translation can introduce milliseconds of reported UTC variation while the underlying monotonic map remains quiet at a much smaller scale.

Our contributions are: (1) a concrete userspace macOS/device architecture that uses captured SOFs as timing evidence while USB control transfers carry the evidence; (2) an integer frame-origin identification procedure constrained by causal marker transactions, separate from rate and phase fitting; (3) a ten-minute case study with complete short-window and held-out prediction populations; and (4) an explicit account of what these measure-

ments cannot establish. The contribution is a method and an auditable experiment, not the invention of USB synchronization, a universal Mac accuracy guarantee, or a gPTP implementation.

The most useful result is the separation of three effects. Short-term monotonic phase variation is sub-microsecond after a declared local linear detrend. Longer-term relative rate variation makes a single ten-minute line a worse model. The UTC translation contributes a separate millisecond-scale effect. This decomposition explains why an attractive early instrument reading does not automatically extend to a longer run or to UTC accuracy.

2 Related work and comparison criteria

2.1 gPTP, PTP and the Mac clock endpoint

IEEE 802.1AS specifies timing and synchronization for time-sensitive networks, using IEEE 1588 mechanisms where applicable. The IEEE working group’s current published list includes 802.1AS-2025; historical evaluations below use their stated earlier editions [2]. gPTP supplies a network timing architecture with synchronization messages, rate relationships and link-delay handling. The USB prototype in this paper does not implement its message set, peer-delay machinery, election or conformance requirements. The conceptual connection is the transport of event timestamps and the need to cross clock domains without treating software delivery delay as event phase.

A particularly relevant implementation is Clark’s *mac-gptp*. It exposes a Mac’s AVB gPTP virtual clock through the private TimeSync framework and can feed a chrony reference-clock socket. Its documented setup uses an AVB-capable Ethernet path and a Linux boundary clock. The README reports roughly 250 ns agreement between successive Mach/gPTP mapping estimates over a few seconds and explicitly retains uncertainty in a fixed adapter-to-Mach cross-timestamp offset [3]. This is closely related to our host-endpoint problem, despite the different transport and API.

The project’s refclock account reports 0.6 μ s RMS over half an hour with tuned chrony settings and approximately 2 μ s excursions during its fastest wander. It identifies oscillator wander by comparing the Mac mapping with a separate adapter-clock mapping and explains the importance of short frequency-estimation histories [4]. These are author-reported system-clock-to-gPTP residuals, not our measurements and not a calibrated UTC benchmark. We did not run that project, configure AVB, or steer the system clock. Its distinct adapter clock also provides an attribution experiment that our USB run lacks: our relative rate changes cannot yet be assigned uniquely to the Mac, USB reference, TC-04 or its upstream source.

Val et al. evaluate an FPGA-based Ethernet/wireless TSN clock-domain translation with an external Calnex measurement. Their Table III reports -4.3 ns mean and 18.7 ns standard deviation for a three-hop Ethernet setup, and -5.5 ns mean and 16.6 ns standard deviation for one 802.11 plus two Ethernet hops [5]. These are externally observed synchronization distributions in their implementation, not residuals after independently fitting short intervals. The numerical similarity of any short-window figure in our study would not establish an equivalent result.

NIST’s calibration study illustrates the missing metrology more directly. Technical Note 2030 reports a 36 ns maximum offset through two transparent clocks with a 6 ns one-sigma uncertainty, and separately discusses source calibration and a power-related anomaly that increased receiver-calibration uncertainty to 100 ns [6]. This is evidence for the importance of reference calibration, not a performance baseline that our experiment can claim to meet.

2.2 USB synchronization and causal delay bounds

Lim, Rosdi and Yap use USB 3.0 Isochronous Timestamp Packets to coordinate device triggering. Their publisher abstract reports approximately 200 ns trigger delay and operation through up to three consecutive hubs [7]. The endpoint is relative triggering among SuperSpeed devices, rather than a full-speed device-to-Mach clock association. It is nevertheless clear prior evidence that USB timing can be much finer than a software transaction interval. Patent literature also describes recovering and synthesizing device clocks from periodic USB structures [8]; we do not present SOF-based synchronization as a new general principle.

WALT’s primary clock-synchronization notes develop causal upper and lower clock-offset bounds from messages in both directions. They report interval widths around 100 μ s on one Linux USB2 path and 800 μ s on a USB3 path [9]. Those are bound widths on named configurations, not standard deviations or physical accuracy floors. We use a related causal principle for *integer identity*, while hardware boundary observations supply the fine phase evidence. Our periodic UTC request intervals deliberately remain wide rather than pretending that a request midpoint is an exact event time.

NTP is a separate network clock-discipline system. RFC 5905 discusses capabilities extending to tens of microseconds on suitable workstations and fast LANs [10]. The Internet pool checks in our experiment do not test that regime. A noisy network observation is not evidence that the underlying NTP-disciplined clock is noisy by the same amount.

3 Clock coordinates and measurement vocabulary

Let C denote the extended device hardware counter and $M_g(C)$ the source-monotonic time supplied by the map generation g applicable to that event. Let O_g be the reported UTC translation and $U_g(C) = M_g(C) + O_g$ the reported POSIX UTC estimate. On the host, let A denote Mach absolute time after tick conversion, K Mach continuous time, and W local wall time. The desired application coordinate would be

$$\hat{U}(K) = M_g\left(\hat{C}\left(\hat{A}(K)\right)\right) + O_g, \quad (1)$$

subject to valid frame identity, map intervals, clock-domain association and source permission. In the tested client this precision UTC service remains withheld. Captured event UTC is available as evidence; it is not a clock served at the current host instant.

We use *detrended phase standard deviation* for the population standard deviation remaining after a declared linear fit. *Drift* is the corresponding fitted fractional rate, in ppm. *Causal innovation* evaluates a new point against a fit containing only earlier observations. *Held-out block prediction* freezes one block’s fit for evaluation on a following block. *UTC accuracy* requires correct epoch/second association and an independently bounded reference and path. These statistics answer different questions. Windowed detrended standard deviation is not an Allan-deviation or time-deviation measurement, and we do not infer a noise process or long-term holdover law from it.


The source-monotonic quantity is recovered from the packet as $M = U - O$ using integer arithmetic before subtraction of the initial anchor and floating-point fitting. This avoids losing fine phase information by first converting epoch-sized nanosecond integers to floating point. Subtracting O changes the coordinate being investigated; it does not repair UTC or remove an error from a delivered UTC service.

4 Implementation: physical capture and evidence transport

4.1 The device counter and SOF capture

The Heltec T096 uses an nRF52840 with a free-running TIMER3 counter. An allocated PPI route connects `USBD.EVENTS_SOF` to `TIMER3.CAPTURE[3]`. Capture channels 0, 1 and 2 retain their existing PPS, radio and current-counter roles. The nominal 16 MHz counter unit is 62.5 ns. Nordic documents that the SOF event is synchronized to the 16 MHz clock before reaching PPI and that the frame counter advances at SOF [11]. Counter resolution and event synchronization are implementation properties, not calibrated bounds on cable-to-counter latency.

The existing owner extends the hardware counter to 64 bits. The SOF age is computed as the unsigned low-32-bit difference between the current counter and CC3; the event coordinate is the current extended counter minus that age. The age must be below 16,000 nominal ticks. A coherent sample combines frame number, latched capture, current extended counter and boot/USB/capture epochs. The read checks that frame and capture values agree across the sampling interval and that the inferred captured boundary is less than one frame old. Races, stale captures and inconsistent increments are rejected. The timer remains free-running; source discipline changes numerical maps rather than resetting or physically tuning the counter.



USB carries both a physical reference and delayed evidence; callback arrival is not phase.

Figure 1: Experimental timing path. TC-04 is a BLE-disciplined reference clock, but the Mac client has no BLE transport. The physical USB event and the message carrying its timestamp have different roles. The independent UTC validation path remains missing.

4.2 The implemented transport is diagnostic EP0

An additional vendor interface coexists with the CDC console. The tested implementation uses endpoint-zero vendor control requests, not a production bulk stream. Kind 1 requests return the most recent captured SOF with source information at a requested 32 Hz. Kind 2 requests return current-counter UTC/source information approximately once per second. Each UTS1 version-1 record is 256 bytes. Separate zero-data OUT markers and 64-byte IN replies support frame identity. The distinction matters: an eventual bulk transport may improve bandwidth and queue behavior, but it would require its own coexistence and latency evaluation.

Records retain the device/boot identity, USB and capture epochs, sequence, frame number, captured and current 64-bit ticks, source/root scope, map generation, anchor, effective/active intervals, original source expiry, rate/phase, age, uncertainty fields and unavailability reasons. Unknown uncertainty uses an explicit sentinel. It is never converted to zero. Serialization and history use bounded storage. The Mac opens the matching vendor interface without seizing the entire device or replacing its configuration.

4.3 Userspace host observations

The Objective-C helper calls `referenceMicroframeWithTime:error:` on an `IOUSBHost` object. Apple’s public documentation describes a recent controller microframe with its associated system time [12]; it does not supply a numerical end-to-end accuracy guarantee. The probe retains the raw returned time, runtime Mach timebase, call brackets, reference age, topology and errors. The observed timebase is 125/3 ns per tick. This conversion is recorded rather than assuming returned integers are nanoseconds.

Interleaved absolute and continuous reads bound their association before and after the call. Their intervals must overlap. This detects an invalid local association; it does not qualify suspend/resume, which was not exercised. Reusing an old absolute-to-continuous offset across sleep is outside the demonstrated method.

5 Identifying a frame without fitting away its identity

5.1 Fine phase is not enough

Let f_i be the device’s 11-bit frame number and k_i its unwrapped full-speed frame ordinal relative to the first accepted sample. Coherent frame differences and counter differences advance k_i ; the host does not assume equal numbering origins. Let (μ_j, h_j) be an Apple reference microframe number and its converted Mach time. With $T_\mu = 125 \mu\text{s}$, define

$$\varepsilon_j = (h_j - h_0) - (\mu_j - \mu_0)T_\mu, \quad (2)$$

$$x_i = 8k_iT_\mu + \varepsilon_{j(i)}. \quad (3)$$

Here x_i is the relative USB/Mach elapsed coordinate used for phase analysis. The host reference need not describe the device request’s arrival instant. Its frame/time pair supplies the controller reference phase.

For a direct full-speed path, candidate device-boundary coordinates take the form

$$b_i(r) = h_{j(i)} + (8k_i + r - \mu_{j(i)})T_\mu, \quad r \in 8\mathbb{Z}. \quad (4)$$

The restriction to one full-speed boundary per eight microframes is a tested host-layout hypothesis for this topology. It is not a rule that one may blindly divide any Apple number by eight or use through an arbitrary hub.

Suppose the wrong origin adds one frame to every b_i . Every phase y_i shifts by the same 1 ms. In a fit $y_i = \alpha + \beta x_i + e_i$, changing α by that amount leaves every e_i unchanged. More generally, any fixed endpoint bias is invisible to the residual after fitting an intercept. Thus smooth residuals alone can establish neither the frame origin nor absolute timestamp accuracy.

5.2 Causal marker bounds

A zero-data OUT request carries a nonzero token. Before completing the request, the device stores an immutable tuple containing the token, frame, captured SOF counter, current counter and epochs. A later IN fetch must match that token and scope. The fetch’s potentially long scheduling delay does not expand the earlier OUT’s causal bracket.

Let $[a_j^-, a_j^+]$ bracket the OUT transaction in Mach time. Let q_j be the counter-derived age of the SOF at the device handler, and $B_j = b_j(0)$. Under the implementation’s capture and timing hypotheses, the origin is constrained by

$$a_j^- - B_j - q_j - \delta \leq rT_\mu \leq a_j^+ - B_j - q_j + \delta. \quad (5)$$

The implementation intersects these bounds across accepted markers and searches the $8T_\mu = 1 \text{ ms}$ grid. It uses a $1 \mu\text{s}$ envelope corresponding to a declared 1000 ppm counter-rate allowance over less than 1 ms. This is a

model assumption, not a calibrated measurement of Apple or PHY bias. A unique grid candidate is therefore an operationally identified frame relationship conditional on those assumptions. It is not an external physical proof of the endpoint offset.

Identification additionally requires a direct full-speed XHCI registry path, at least eight frame-layout checks with no failures, and at least sixteen accepted markers. The host compares its current-microframe and frame-number APIs to test the omitted-microframe layout [13]. Boot/USB/capture changes revoke old identity. Isolated bad tuples remain rejected; repeated failures or a sufficiently large gap fence the model instead of extending an old association indefinitely. Hubs, reset, suspend and alternate controller families require separate validation.

6 Causal mappings and source permission

The source map applicable to the event is part of the measurement. Device publication retains two immutable map intervals, closes a prior interval at activation and opens its successor only when publication is complete. A capture in a publication gap is unavailable. A later generation cannot retrospectively authorize an earlier event. Original source expiry clips permission even while USB is connected.

The counter phase is $y_i^C = 62.5 \text{ ns} (C_i - C_0) - x_i$. For valid source observations, using a source-local origin x_{s0} ,

$$y_i^M = (M_{gi}(C_i) - M_0) - (x_i - x_{s0}), \quad (6)$$

$$y_i^O = O_{gi} - O_0, \quad (7)$$

$$y_i^U = y_i^M + y_i^O. \quad (8)$$

This identity isolates UTC translation from monotonic transfer. All three series use the same captured USB-boundary coordinate. The millisecond effect in y^U therefore cannot be explained solely by USB callback scheduling.

The live counter-to-Mach model uses a bounded history of up to 1024 accepted observations, approximately 32 s at the requested rate. A new observation is first evaluated against the *previous* fit, once at least 64 points exist; only then is it added. This yields causal counter-map innovations. It is distinct from a full-run retrospective fit and from the source prediction experiment below. Future projection is limited to the implementation’s finite horizon; a good historical fit does not imply unlimited holdover.

Source validity is separate from connection, frame identity and display freshness. An invalid source record clears captured-UTC display permission. A conservative host request-start deadline and the original device permission limit freshness. The instrument can continue measuring the raw counter after source time becomes unavailable, but it must label that change of measurement population. Its precision UTC shadow remains gated by unresolved UTC association, even after frame identity is operationally resolved.

7 Experimental design and analysis

The selected experiment ran on 15 October 2026 from 17:06:12 to 17:16:12 UTC, for 600.012660 s by the host continuous clock. TC-04 ran the installed T096 application 0.15.45, with the frozen macOS ARM64 client 0.11 and its packaged helper. The OS reported macOS 26.7, build 25G229. Registry evidence identified a direct full-speed device beneath Apple’s XHCI controller. The precise marketed Mac model was not established in the capture; this is one Apple silicon configuration, not a sample of the M-series family [14].

Bounded read-only console checks immediately before and after the run verified TC-04’s actual selected parent as TC-01, unchanged boot, locked state and valid map. Valid USB source packets retained one root/root-boot/source-epoch/path scope. These observations support the recorded direct path but do not substitute for an uninterrupted independent RF trace. No device was flashed, reset or repinned for this capture. The Mac’s system

clock was not adjusted. Background desktop work was not a controlled load treatment, and temperature and power were not measured.

The helper requested precision observations at 32 Hz and periodic UTC at approximately 1 Hz. All raw records, errors and unavailable responses were retained. A frozen binary replayed the capture forward, and an independent centered NumPy least-squares analysis reproduced the reported phase statistics. Original evidence and analysis were sealed and verified against local mirrors; a same-host mirror is an integrity check, not a geographically independent backup.

For a population $\mathcal{W}$ of N valid observations in a declared window, we fit $\hat{y}_i = \hat{\alpha} + \hat{\beta}x_i$ and report

$$\sigma_{\mathcal{W}} = \sqrt{\frac{1}{N} \sum_{i \in \mathcal{W}} (y_i - \hat{y}_i)^2}, \quad d_{\mathcal{W}} = 10^6 \hat{\beta} \text{ ppm.} \quad (9)$$

The slope uses like time units. Windows are assigned by recorded host capture elapsed time; regression uses the captured event coordinate. Every complete disjoint window of 10, 26, 30 and 60 s is included. The 26 s analysis covers 0–598 s; the others cover 0–600 s. No source gap is interpolated. The statistics are conditional on accepted observations and may underrepresent behavior during missing intervals.

For held-out prediction, the model fitted to block j is frozen and evaluated on the next disjoint block $j + 1$ of the same duration. The evaluation block contributes neither fit parameters nor bias removal. Signed errors are pooled and their uncentered RMS, absolute P99 and maximum reported. The first block trains only. This is an offline causal replay test, not a live system-clock discipline or prospective independent-run validation. Window lengths were selected during post-run exploration; the strongest next experiment will fix them before acquisition.

The 18,332 points are serially correlated observations from one run. Counts of windows are not counts of independent physical trials. We do not attach IID confidence intervals, infer an oscillator noise class, or treat the standard deviation of fitted phase as an uncertainty certificate.

8 Results

8.1 API access, frame identity and availability

All 19,191 helper samples returned a valid reference pair and overlapping absolute/continuous brackets. Median API call duration was 13.000 μ s, P99 30.541 μ s and maximum 301.875 μ s. Reference age at return had median 65.750 μ s, P99 125.504 μ s and maximum 222.958 μ s. These are call/reference-age diagnostics, not errors in the associated timestamp. Precision request errors were zero.

The bridge first reported its unique frame origin at the 4.085 s status observation and remained identified afterward. It retained one hardware segment, 2386 markers and 2399 host-layout checks with zero layout failures. Thirteen markers and 98 boundary tuples were rejected. The raw observation layer counted 19,125 SOF-valid tuples; the stricter bridge accepted 19,093. The different denominators describe different checks and are not interchangeable. The retained frame stream included 292 observed 11-bit wraps.

Of the 830 unavailable precision-source records, 617 report publication busy, 152 have no applicable permitted map, and 61 have invalid/incoherent capture. Valid source packets constitute 95.675% of precision requests; this is not a continuous-service percentage. The largest interval between requests carrying valid precision-source records was 95.162 ms. The longest point-sampled unavailable status interval was 1.027 s. All eight unavailable periodic UTC responses report publication busy. TC-04 was still locked to its direct parent with advancing accepted corrections after the experiment; device lock does not itself authorize every historical USB event.

Table 1: Observed populations. Missing observations are retained in the capture and excluded only from the metric requiring their unavailable coordinate. Point-sampled availability is not continuous coverage.

Population	Valid/accepted	Unavailable/rejected	Meaning
Precision source records	18,361	830	Of 19,191 requests
Bridge boundary tuples	19,093	98	Counter/reference coherence
Source phase in bridge	18,332	—	Intersection of source and bridge validity
Periodic UTC records	577	8	Of 585 observations
Source-valid status	568	26	Of 594 status observations
NTP queries	10	0	Read-only network checks

8.2 Ten-minute phase and shorter windows

Table 2 and Fig. 2 show why a single jitter number is insufficient. Source-monotonic phase advances 2.794 ms from first to last accepted event, mainly due to a relative rate of +4.635 ppm. After one ten-minute fit, the remaining phase has 3.956 μ s standard deviation. The raw counter’s corresponding values are +5.475 ppm and 4.622 μ s. These are relative rates of composite clock paths, not proof of which oscillator wandered.

Table 2: Full-run phase statistics against USB/Mach. Residual statistics remove one retrospective line over the whole population; UTC and monotonic coordinates remain separate.

Measurement	Result	Population / interpretation
Source monotonic phase σ	3.956 μ s	18,332 points; drift +4.635178 ppm
Source residual P95/P99/max	6.678/8.091/8.374 μ s	Absolute residuals of the same fit
Raw source phase σ	797.831 μ s	Drift retained
Counter phase σ	4.622 μ s	19,093 points; drift +5.474520 ppm
Causal counter innovation σ	0.270 μ s	19,029 prior-model predictions; raw innovation SD
Reported UTC phase σ	1.465 ms	18,332 captured events
UTC-translation phase σ	1.465 ms	Same events; 7.984 ms peak-to-peak

Table 3: All complete disjoint windows of each duration, using a separate retrospective linear fit within each window. Values are source-monotonic phase, not UTC accuracy.

Window (s)	Windows	Median σ (μ s)	Worst σ (μ s)	Drift range (ppm)
10	60	0.065	0.139	+4.442 to +4.801
26	23	0.117	0.176	+4.468 to +4.786
30	20	0.155	0.288	+4.472 to +4.758
60	10	0.279	0.593	+4.524 to +4.756

The first 26 s of this capture gives 0.158 μ s standard deviation, consistent with the earlier operator screenshot’s 0.18 μ s after 26 s. Across all 23 complete 26 s blocks, the range is 0.061–0.176 μ s. Thus the longer-run result does not imply that short-term stability became twenty times worse. It describes a longer horizon under a more restrictive single-line model. The ten 60 s fitted rates vary by approximately 0.232 ppm. Their standard deviations remain 0.179–0.593 μ s.

An earlier 600 s development capture contained only 248 valid source observations, approximately 8 s, despite having an extensive raw counter stream. Its attractive 0.062 μ s source statistic was not a ten-minute source result.

This historical failure motivated retaining observation counts and validity denominators with every readout. It is not pooled into the present run.

8.3 Held-out phase prediction

The previous-block test in Table 4 avoids using the evaluated block to choose its own slope. At 30 s the pooled RMS is $1.053\ \mu\text{s}$, absolute P99 $3.285\ \mu\text{s}$, and maximum $3.662\ \mu\text{s}$. At 60 s, RMS rises to $3.359\ \mu\text{s}$ with a $10.329\ \mu\text{s}$ maximum. The difference from the retrospective window residual is expected: changing relative frequency is not known in advance.

Table 4: Offline held-out prediction using only the preceding disjoint block. Each model is frozen for the following equal-duration block. RMS retains signed bias; no evaluation-block recentering or refitting.

Fit / test (s)	Test points	RMS (μs)	Abs. P99 (μs)	Abs. max (μs)
10 / 10	18,035	0.338	1.111	1.535
26 / 26	17,501	0.973	3.103	3.930
30 / 30	17,439	1.053	3.285	3.662
60 / 60	16,544	3.359	9.645	10.329

8.4 UTC translation, local wall time and NTP

Equation 8 identifies the dominant remaining effect: the detrended UTC-translation series has $1.465443\ \text{ms}$ standard deviation, while reported UTC phase has $1.464967\ \text{ms}$. Their near equality supports attribution of the *reported-coordinate variation* to that translation. It does not identify its upstream cause or independently measure the actual UTC error. The translation spans $7.983543\ \text{ms}$ peak-to-peak. Improving USB scheduling alone cannot remove that term.

The 577 valid periodic observations keep local-wall reads bracketing the complete current-counter USB request. For captured UTC u_i and ordered wall reads $[w_i^-, w_i^+]$, the measured offset interval is

$$u_i - w_i^+ \leq U - W \leq u_i - w_i^-, \quad (10)$$

conditional on the event lying in that request and wall time remaining monotonic over the bracket. Bracket widths range $1.991\text{--}3.148\ \text{ms}$. Their mean offset interval, TC-04 minus local wall, is $[-25.261, -22.390]\ \text{ms}$. The corresponding observed RMS bounds are $[22.441, 25.305]\ \text{ms}$, with largest absolute upper bound $29.174\ \text{ms}$. These are sampled request-interval diagnostics; a midpoint standard deviation of $1.508\ \text{ms}$ includes USB scheduling uncertainty and is not the fine boundary-based stability measure.

Ten read-only SNTP queries to `pool.ntp.org` succeeded. Server-minus-local offsets range $+22.410$ to $+29.325\ \text{ms}$, with reported uncertainties ± 52.884 to $\pm 96.552\ \text{ms}$. DNS selected four server addresses. The raw spread of those ten offsets is $1.952\ \text{ms}$ standard deviation, combining network path and server selection effects. It cannot establish NTP oscillator jitter or a stability ratio against our source. The OS network-time enabled state remained unverified because the configuration tool required administrator access. No query set or slewed the clock.

9 What can be compared with prior work?

The USB experiment demonstrates an accessible counter-to-host observation path on hardware without an AVB Ethernet clock. It does not demonstrate that USB has a better ultimate timing limit than Ethernet, that its accuracy matches hardware PTP, or that it outperforms a well-configured NTP installation. Its strongest comparison is




Figure 2: One capture, three different questions. (a) Relative source phase retains its accumulated drift. (b) A single ten-minute line removes the mean rate but leaves slower relative rate variation. (c) UTC retains a much larger variation even at the same captured boundaries. PGFPlots displays actual first/last and extrema samples within each one-second bin; all accepted points contribute to the statistics. Missing records are not filled or treated as zero.




Figure 3: Relative source phase error predicted from the preceding 30 s block, with no test-block fit. Each new 30 s block uses a new model trained on its immediate predecessor. The first 30 s supplies training only. Points are visually reduced by one-second extrema; RMS and quantiles use all 17,439 evaluation points. This is an offline timing-coordinate prediction, not a live served UTC clock.

Table 5: Comparators retain their original endpoints and metric meanings. Numbers are descriptive context, not a ranking; none was measured head-to-head in this study.

Work	Selected reported result	Endpoint and essential distinction
This USB/Mac study	0.155 μ s median 30 s σ ; 1.053 μ s next-block RMS	Captured source-monotonic coordinate versus USB/Mach; retrospective fit and held-out replay are separate. Absolute bias unknown.
mac-gptp [3, 4]	Approximately 250 ns mapping agreement; 0.6 μ s ref-clock RMS	Author-reported AVB/Mach mapping consistency and chrony-driven wall-clock residual, respectively. Private API; fixed cross-timestamp uncertainty retained.
Lim et al. [7]	Approximately 200 ns trigger delay	USB 3.0 ITP-based device triggering; different bus speed, endpoint and multi-hub experiment.
WALT [9]	Approximately 100–800 μ s selected USB offset-bound widths	Causal software message bounds on named Linux paths; not a phase σ .
Val et al. [5]	18.7 ns SD, three-hop Ethernet; 16.6 ns SD, hybrid path	External synchronization measurement of FPGA TSN architectures; no local per-window line subtraction equivalent to ours.
NIST TN 2030 [6]	36 ns maximum offset, 6 ns one-sigma uncertainty through two TCs	Calibrated PTP timing chain; source and reference uncertainties investigated separately.

methodological: physical event timestamps and explicit clock-domain mappings remove a class of delivery-time ambiguity that a callback-only measurement retains.

The macOS gPTP case also prevents an overly broad novelty claim. Mapping an external time domain to Mach time already has a relevant implementation precedent. Our engineering contribution is the public IOUSBHost/SOF path, its frame-origin constraints and source-map evidence handling on this device. A controlled experiment using the same Mac, source, comparison endpoint and independent reference would be required to compare gPTP and USB numerically. No such experiment is reported here.

10 Uncertainty, failure cases and limits

10.1 A partial budget is not an accuracy certificate

An application UTC error can be decomposed schematically as

$$e_U = e_{\text{source}} + e_{\text{map}} + e_{\text{capture}} + e_{\text{identity}} + e_{\text{host}} + e_{\text{topology}} + e_{\text{holdover}} + e_{\text{scale}}. \quad (11)$$

This is a dependency ledger, not a fitted error model. The observed standard deviations are neither deterministic bounds on its terms nor evidence that they are independent. If each term had a justified absolute bound B_i , a conservative sum would bound the total; an RSS combination would require an appropriate probabilistic and covariance model. Several terms currently have no calibrated bound, so no finite sub-100 μ s total budget is assigned.

The 62.5 ns timer unit does not bound SOF detection or PPI latency. The marker’s 1 μ s allowance does not calibrate Apple’s reference. The API call duration does not measure timestamp error. The unique integer candidate does not remove constant cable/controller/PHY bias. The observed short-window phase scatter combines the device, source mapping and host reference; common-mode or constant errors can remain invisible. A neighboring reference clock or calibrated physical association to Mach is needed to separate them.

10.2 Availability and inference limits

Publication-busy responses can be correct refusals: a software map transition must not quietly retimestamp an older event with a newer generation. Their frequency is still a service limitation. The 152 no-map responses and 61 invalid captures need causal tracing through map activation, expiry and coherent sampling. It would be unsound to improve a coverage percentage by relaxing those guards or imputing time into missing intervals.

Only one direct full-speed XHCI path, one source scope and one ten-minute run are represented. There is no controlled CPU/I/O stress population, thermal treatment, sleep/reconnect trial, alternate-Mac cohort or hub-depth sweep. The beginning and end parent checks do not observe every intervening RF state. A phase fit conditional on available samples does not establish continuous precision service. The 30 s window choice and comparison were exploratory; a future predeclared repeat is needed to estimate repeatability and selection effects.

The periodic wall-time bracket is valid only for its captured instant and its stated assumptions. Sparse NTP results are not propagated as a guaranteed time map. Local wall time showed a fitted rate of approximately +3.045ppm relative to Mach during this run, but the recording does not establish the OS discipline state or the cause of that rate. POSIX UTC estimates also require an explicit policy for leap-second boundaries before production use; no such boundary was exercised.

10.3 Trust and resource boundaries

The Mac binary contains no BLE client path and uses one selected USB interface, while the RC continues upstream radio service. Counter, capture-channel, PPI and endpoint coexistence therefore remain real resource constraints. This paper reports the measured configuration rather than authorizing a firmware change or a broader independent test suite. Exact USB serial/interface selection and local exclusive ownership are operational protections, not cryptographic authentication against a malicious USB device. The trusted host/device experiment does not establish adversarial time integrity.

11 A validation path toward an accuracy claim

The next experiment should specify the claim before collecting data: for example, a bound or percentile of a *served* source-monotonic coordinate at an independently timed host event, with an explicit availability denominator. That claim is different from absolute UTC, which additionally needs a qualified source and second association. Success criteria should retain startup, source loss, rejected events and all recovery tails.

First, resolve the UTC translation’s millisecond variation. Record the upstream physical event, source UTC/second association, original issuance and each map activation alongside the exported translation. Determine whether the variation arises in source association, publication, upstream correction or a different mechanism. The present algebra localizes the varying field but does not establish its root cause.

Second, validate the integer frame association and endpoint bias with independent evidence. Decoded USB SOFs plus a hardware trace can test the device capture and any frame slip. An analyzer alone does not automatically associate that event with Mach time: a separately justified Mac-clock cross-association is also required. Repeat the identity procedure across frame wrap, USB reset, reconnect, suspend/resume and each proposed topology, retaining ambiguous or empty candidate sets.

Third, predeclare a causal estimator and repeat its evaluation. Fix the training window, prediction horizon, minimum sample/coverage policy, segmentation and outlier rules before a new run. Include longer holds, deliberate consumer stalls and controlled host/device load. A good delivered clock must preserve original source expiry and withdraw permission predictably, not merely retain a low residual on surviving points.

Finally, compare USB, gPTP and local/NTP timing at a common endpoint against the same independent reference where the hardware permits. Retain signed bias, RMS, percentile/maximum error, observation gaps and

source uncertainty. Keep fixed corrections trained on one population separate from validation on another. These steps would convert a promising stability demonstration into a falsifiable end-to-end timing claim. They are proposed work, not completed evidence.

12 Discussion and conclusion

USB can supply two useful products at once: timestamped evidence from an attached device and a physical reference with which to cross into a desktop clock domain. Their usefulness does not depend on making application callbacks deterministic. Device-event correlation, recorded sensor fusion and short-lived shadow coordinates can benefit from this separation even before a global wall-clock service is justified. Each application must still enforce the mapping’s scope and lifetime.

The experiment demonstrates short-term relative stability across an entire ten-minute capture: every complete 30 s window has detrended source phase standard deviation below $0.288\,\mu\text{s}$, with a median of $0.155\,\mu\text{s}$. Predicting the following 30 s block using only the preceding block gives $1.053\,\mu\text{s}$ RMS and $3.662\,\mu\text{s}$ maximum relative phase error. A single ten-minute line leaves $3.956\,\mu\text{s}$ standard deviation, showing why a stability figure must state its horizon and fitting rule. These are measured coordinate relationships on one composite source/device/host path.

The unresolved result is equally important. Reported UTC retains millisecond variation, and absolute host-reference bias remains unknown. Neither a 1 ms frame interval nor a $0.155\,\mu\text{s}$ residual determines absolute accuracy. A credible USB time-transfer service needs correct event identity, an applicable causal source map, a host-clock association and a qualified UTC source. This preprint establishes an experimental path and its measurement vocabulary while keeping that final qualification open.

Data, code and reproducibility

This is a local review draft, not a submitted or published paper. The companion `paper/usb-macos-evidence-20261015.json` records measured populations, frozen input hashes, literature provenance and uncertainty status. `paper/figures/usb-macos-20261015/derive.py` verifies the phase CSV against its parent seal, regenerates inline PGFPlots coordinates and cross-checks the full-run statistic. The standalone source embeds its figures and bibliography and needs no private raw payload to compile. A relative-phase review CSV omits physical identities and epoch-sized UTC values; it is a derived subset, not the complete availability record.

The two immutable local evidence archives are indexed by `TC04_DIRECT_USB_600S_20261015_01` and `TC04_SHORT_WINDOW_STABILITY_20261015_01`. Their raw/private records remain outside website assets and any public dataset. A public code/data release, durable off-host preservation and licensing review remain pending. The research task performed only read-only evidence analysis and literature retrieval; it did not access hardware, change clock configuration or update the website.

Author contributions and assistance disclosure

Jonathan Yantis supplied the hardware, original protocol/firmware, experimental direction and research decisions. OpenAI Codex assisted implementation and experimental tooling in the project, literature research, source/evidence review, analysis, figure/table generation, writing and LaTeX preparation for this draft. No AI system is listed as a human coauthor. The human author retains responsibility for reviewing evidence, claims, references, data privacy and the requirements of any eventual venue. No affiliation, acceptance, DOI or arXiv identifier is claimed.

A Reproduction notes and guard contract

The main capture used the retained packaged client with `-duration 600 -rate 32 -no-interactive` and a fresh JSONL path. The frozen client replay generated boundary-phase CSV without extending source permission. The independent short-window analysis verifies that its 18,332 source-event coordinates increase strictly and share one scope. It includes all complete blocks, checks that every evaluated block follows its training data and retains signed error in held-out RMS. The exported statistics use full-resolution data; the plots use one-second extrema only for display.

UTS1 kind 1 is requested through vendor IN request 0x53; kind 2 through 0x52. Both retain the source map generation and original expiry. The marker uses OUT 0x54 with a token and later IN 0x55 for the immutable 64-byte tuple. The tested interface is vendor class/subclass/protocol FF/54/01. This is an experimental transport contract, not a standard USB time class. A production bulk interface was not tested.

Boundary acceptance checks include scope consistency, strictly advancing sequence/counter, frame increment consistent with the nominal counter rate, a capture age below one frame, valid Apple reference and overlapping absolute/continuous association brackets. An isolated inconsistent tuple is rejected while retaining the last good boundary; repeated errors or a large gap fence the model. Validity at an event is evaluated against its original map interval and current source expiry. Unavailable data is never reconstructed from a later map.

The full raw capture SHA-256 is `be3e0ff1fd217b2092c17cb1e83255cd5816fcfa9c30ee3daa338f7124756bfd`. The bridge phase CSV SHA-256 is `fc493b959aaa4b2f99d7ea9c8443338b7f45d2f604ea7bdf15b222d2f94020e7`. The selected firmware source receipt identifies commit `68b467a308e0462d7dd0df41678122c80b00e8b0`; this research run verified identity/version/boot before and after but did not repeat an image readback. Companion manifests distinguish the earlier installation evidence from this capture’s read-only checks.

References

- [1] J. Yantis with disclosed OpenAI Codex assistance, *Passive Precision Time over BLE: V14 Boundary Clocks, Causal Timestamping, and Prospective One-Relay Validation*, local review preprint, 12 October 2026. Repository paper/ptob-v14-preprint.tex. Prior radio evidence is not pooled into this USB result.
- [2] IEEE 802.1 Working Group, *IEEE Std 802.1AS-2025: Timing and Synchronization for Time-Sensitive Applications*, published-standard listing and maintenance project scope. <https://1.ieee802.org/tsn/>; <https://1.ieee802.org/maintenance/802-1as-2025-revision-timing-and-synchronization-for-time-sensitive-applications/>. Accessed 15 October 2026. Public project material inspected; no claim of a clause-by-clause conformance review.
- [3] J. Clark, *mac-gptp: Use gPTP on macOS for timing*, README and implementation documentation, commit `e133922d72282e1a41835a19e575663d292a21c5`, 4 September 2026. <https://github.com/jclark/mac-gptp/tree/e133922d72282e1a41835a19e575663d292a21c5>. Primary project report, accessed 15 October 2026; not independently reproduced here.
- [4] J. Clark, *gptp-refclock*, technical notes at the same pinned revision. <https://github.com/jclark/mac-gptp/blob/e133922d72282e1a41835a19e575663d292a21c5/docs/gptp-refclock.md>. The reported 0.6 μ s RMS is a system-clock residual against the project’s gPTP reference, not calibrated absolute UTC error.
- [5] I. Val, Ó. Seijo, R. Torrego and A. Astarloa, “IEEE 802.1AS Clock Synchronization Performance Evaluation of an Integrated Wired-Wireless TSN Architecture,” *IEEE Transactions on Industrial Informatics*, 2022, doi:10.1109/TII.2021.3106568. Author manuscript arXiv:2109.09409, 2021, Table III. <https://arxiv.org/abs/2109.09409>.
- [6] M. A. Weiss, Y. Li-Baboud, D. Anand, K. G. Brady, P. A. Boynton, C. T. Nguyen, M. J. Burns and A. M. Gopstein, *A Calibration of Timing Accuracy in the NIST Cyber-Physical Systems Testbed*, NIST Technical Note 2030, December

- 2018, doi:10.6028/NIST.TN.2030. <https://www.nist.gov/publications/calibration-timing-accuracy-nist-cyber-physical-systems-testbed-0>. The Technical Note summary, rather than the differing conference abstract, supplies the selected maximum/uncertainty wording.
- [7] C. H. Lim, B. A. bin Rosdi and C. F. Yap, “Synchronization of multiple USB 3.0 devices using Isochronous Timestamp Packet,” *Computer Standards & Interfaces*, 2017, doi:10.1016/j.csi.2016.06.006. <https://www.sciencedirect.com/science/article/pii/S0920548916300423>. Publisher abstract/highlights inspected; full experimental article not available in this review.
- [8] P. G. Foster, *USB based synchronization and timing system*, US8943351B2, 2015. <https://patents.google.com/patent/US8943351B2/en>. Primary disclosure of USB periodic-signal clock synchronization; a patent disclosure is not an independently reproduced benchmark.
- [9] Google WALT contributors, “Clock Synchronization,” primary implementation notes preserved in Android Open Source Project. <https://android.googlesource.com/platform/external/walt/+/HEAD/android/WALT/app/src/main/jni/README.md>. Accessed 15 October 2026; named-machine interval widths retained as reported.
- [10] D. Mills, J. Martin, J. Burbank and W. Kasch, “Network Time Protocol Version 4: Protocol and Algorithms Specification,” RFC 5905, June 2010. <https://www.rfc-editor.org/rfc/rfc5905.html>.
- [11] Nordic Semiconductor, *nRF52840 Product Specification*, v1.2, USB isochronous transactions / SOF event and TIMER capture descriptions. https://docs-be.nordicsemi.com/bundle/nRF52840_PS_v1.2/raw/resource/enus/nRF52840_PS_v1.2.pdf. Accessed 15 October 2026; source documents 16 MHz synchronization before PPI, not a calibrated installed-device timing bias.
- [12] Apple Inc., “referenceMicroframeWithTime:error:,” IOUSBHostObject documentation and local macOS SDK declaration. <https://developer.apple.com/documentation/iousbhost/iousbhostobject/referencemicroframewithtime:error:>. Accessed 15 October 2026. The public API is distinguished from DriverKit interfaces and private TimeSync.
- [13] Apple Inc., “frameNumberWithTime:,” IOUSBHostObject documentation and local SDK declaration, accessed 15 October 2026. <https://developer.apple.com/documentation/iousbhost/iousbhostobject/frameNumberWithTime:>. The inspected SDK explicitly describes omission of the microframe component.
- [14] J. Yantis with OpenAI Codex assistance, *TC-04 direct TC-01 USB comparison* and *TC-04 shorter-window stability results*, sealed local evidence, 15 October 2026. Companion paper/usb-macos-evidence-20261015.json; literature retrieval and claim audit in paper/provenance/usb-macos-literature-20261015.json. No public raw-data release is claimed.