High-Precision Time Synchronization and Autonomous Maintenance for LEO Satellite Constellations Based on High-Stability Crystal Oscillators

Authors: L. MU[1,2], X. HU[1], M. WU[1] and J. LI[1,2]
[1] Shanghai Astronomical Observatory, Chinese Academy of Sciences, Shanghai, China; [2] School of Astronomy and Space Science, UCAS, Beijing, China
Contact: mulei@shao.ac.cn

IGS 2026 WORKSHOP
SANTIAGO DE CHILE
33°26' S 70°39' O

INTRODUCTION & OBJECTIVE

- Large-scale LEO constellations require nanosecond-level timing for beam switching, TDMA scheduling, high-rate modulation, cooperative sensing, inter-satellite networking, and distributed PNT services.
- Conventional timing architectures rely on continuous GNSS reception or onboard atomic clocks. GNSS timing is vulnerable to interference and blockage, while atomic clocks add payload cost, power consumption, volume, and system complexity.
- A high-stability OCXO is compact and low power, but its frequency offset, temperature sensitivity, and aging drift cause time error to accumulate quickly when it is left free-running.
- The key idea is to convert short visibility windows into periodic correction opportunities: the satellite locks to an external reference only during visible arcs, then maintains time autonomously during long non-visible intervals.
- This approach emphasizes physical signal generation: the onboard system directly produces a 1PPS reference instead of only estimating clock offset in a backend processing system.

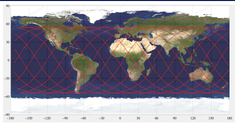

Constellation ground-track distribution used to derive intermittent synchronization opportunities.

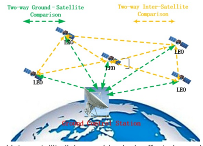

Two-way star-ground and inter-satellite links provide clock-offset observables during visible arcs.

Objective and Contributions

- Adopt a realistic 7 min closed-loop synchronization + 113 min open-loop holdover cycle for an 800 km Walker-Delta LEO scenario.
- Integrate two-way star-ground / inter-satellite measurements, temperature feedforward α(T), Kalman estimation, LQR voltage control, and aging prediction.
- Use measured hardware parameters—voltage sensitivity κ, TMP117 thermal readings, DAC resolution, and aging drift—to drive a real OCXO control board.
- Revised target and reported statistics use STD / 1σ: <5 ns with peak-to-peak <30 ns while continuously outputting a local 1PPS signal.

METHODS

Constellation and Timing Links

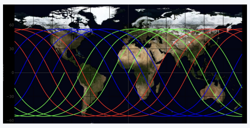

Representative tracks from multiple planes show clustered visibility rather than continuous access.

- Walker-Delta constellation: 36 satellites, 6 planes, 6 satellites per plane, 800 km altitude, and 55° inclination for representative China-region service.
- Orbital period is approximately 100–101 min. For a Beijing-region ground station, a single visible pass is about 7.5 min, leaving most of the cycle for autonomous holdover.
- Two-way time transfer reduces first-order path-delay sensitivity and estimates the star-ground or inter-satellite clock offset, but residual OD/geometry effects and terminal-delay calibration uncertainty still enter the system error budget.
- Inter-satellite links can relay the timing state from satellites with ground access to satellites currently outside the ground-station visibility window.
- The adopted a + b schedule intentionally represents a constrained measurement pattern: short measurement availability followed by long local maintenance.

State Estimation and Control

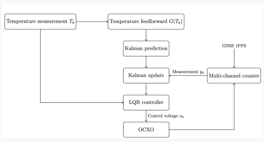

Temperature feedforward, Kalman filtering, and LQR optimal control form the core control loop.

- State vector x = [x0, x1, x2]^T represents time bias, frequency offset, and frequency drift rate; these terms correspond to phase, frequency, and aging-related dynamics.
- The state transition includes the control-voltage term κu_k and a temperature feedforward term α(T_k), where α(T) is identified using piecewise temperature-frequency calibration.
- Only x0 is directly observed from 1PPS time-offset measurements; Kalman filtering reconstructs x1 and x2 while balancing measurement noise and process noise.
- LQR computes a smooth DAC tuning voltage. Increasing phase weight accelerates time-bias removal, while frequency/drift weights suppress fluctuations and slow aging trends.
- During holdover, the last estimated state, real-time temperature, and aging model are propagated forward until the next synchronization opportunity becomes available.

Hardware Validation Platform

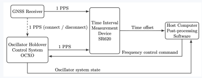

Ground-based platform emulates intermittent timing-link availability and autonomous holdover.

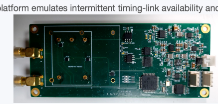

OCXO voltage-control board with AD5791 DAC and four TMP117 temperature sensors.

- GNSS-disciplined 1PPS is used as the reference; an SR620 time-interval counter measures the offset between the OCXO-generated 1PPS and reference 1PPS once per second.
- The AD5791 DAC provides microvolt-level voltage tuning with theoretical LSB ≈ 1.955 μV, enabling fine control around the OCXO EFC input.
- Four TMP117 chips are installed near the OCXO. They provide 0.0078 °C resolution and typical ±0.1 °C accuracy, supporting second-by-second temperature feedforward.
- The STM32F103 MCU reads TMP117 sensors via I²C, drives the DAC via SPI, and exchanges voltage/temperature/status data with the host computer through RS-232 at 1 Hz.
- Voltage-to-frequency sensitivity κ, temperature coefficients α(T), and aging drift (~0.232 ppb/day) are experimentally characterized and embedded into the model.

RESULTS

Statistics use STD / 1σ

Condition	STD / 1σ	Peak-to-peak	Interpretation
Free-running 24 h	-	119.9 μs	Raw drift without control
Drift-removed 24 h	-	491 ns	Residual after first-order removal
Controlled 2 h cycle	1.1399 ns	4.156 ns	7 min closed-loop + 113 min holdover
Controlled 24 h	3.931 ns	17.386 ns	12 repeated synchronization cycles

One-week autonomous maintenance

4.89 ns
STD / 1σ

28.69 ns
peak-to-peak

Performance remains below the <50 ns resilient PNT threshold while validating stable long-duration holdover.

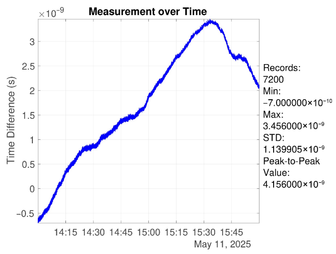

2 h cycle: STD/1σ = 1.1399 ns; peak-to-peak = 4.156 ns.

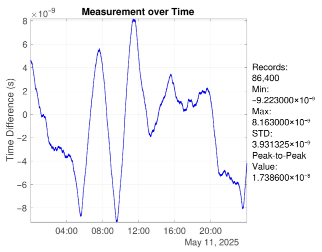

24 h operation: STD/1σ = 3.931 ns; peak-to-peak = 17.386 ns.

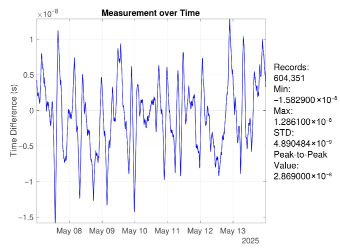

One-week operation: STD/1σ = 4.89 ns; peak-to-peak = 28.69 ns.

- The free-running OCXO reaches 119.9 μs peak-to-peak over 24 h; after removing first-order frequency drift, the residual is still 491 ns, confirming that raw OCXO holdover alone is insufficient.
- Under the proposed control strategy, the time-domain error remains at the nanosecond level over 2 h, 24 h, and one week, demonstrating that the closed-loop/open-loop schedule can suppress cumulative drift.
- Frequency-domain analysis in the paper shows long-term stability at τ = 1000 s improves from 3.19×10^{-12} to 6.54×10^{-13} under control, while maintaining practical short-term behavior.
- The result supports the central claim: intermittent time transfer plus model-based local compensation can bridge low-cost oscillator hardware and stringent constellation timing requirements.

CONCLUSIONS

- A low-cost high-stability OCXO can support nanosecond-level autonomous timekeeping when combined with intermittent two-way timing links, temperature feedforward, Kalman filtering, LQR control, and aging compensation.
- The 7 min closed-loop + 113 min open-loop schedule converts limited LEO visibility into a practical synchronization-and-holdover cycle, reducing dependence on continuous GNSS or onboard atomic clocks.
- The hardware implementation is essential: AD5791-based microvolt tuning, TMP117 temperature sensing, and RS-232/SPI/I²C embedded control allow the algorithm to act on a real physical oscillator.
- Using revised STD/1σ statistics, demonstrated performance is 1.1399 ns over 2 h, 3.931 ns over 24 h, and 4.89 ns over one week; one-week peak-to-peak error is 28.69 ns.
- The architecture directly outputs a local physical 1PPS signal, making it suitable for decentralized, resilient, lightweight constellation timing in weak-GNSS or GNSS-denied scenarios.
- Future work should extend the validation to multi-node timing links, orbit-determination residuals, terminal-delay calibration, and onboard environmental perturbations.

Implication

- The experiment indicates that the dominant OCXO error sources can be transformed from uncontrolled drift into predictable and controllable model terms.
- For constellation deployment, the same framework can be extended from one satellite-clock node to multi-satellite timing dissemination, where ISL contacts reduce the effective holdover duration.
- The measured nanosecond-level performance should be interpreted as method-induced timing performance under the tested measurement model and control schedule, not a complete end-to-end in-orbit error budget.

Scan Full Paper

DOI: 10.3390/s26061839

POSTER ID
P1 06