ORNL/SPR-2026/4615

CAST — Center for Alternative Synchronization and Timing —

Technical Bulletin #013:

Summary: While Global Navigation Satellite systems (GNSS) offer a very accurate and stable source of time and synchronization they also suffer from problems such as jamming, spoofing, and malfunctions. Two Way Satellite Time and Frequency Transfer (TWSTFT) is one of the few techniques that can rival the synchronization accuracy provided by GNSS. In this bulletin, the accuracy of a TWSTFT network consisting of several geographically disparate sites is explored and quantified.

TWSTFT Network Testing

Introduction

Over the past several decades, billions of dollars have been spent putting extremely precise atomic clocks into space [1]. These atomic clocks orbit Earth onboard dozens of satellites and altogether form several independent Global Navigation Satellite System (GNSS) constellations. A GNSS constellation, such as the Global Positioning System (GPS), allows users anywhere in the world to synchronize their time to the much more accurate time signal being transmitted by each satellite in that constellation. Additionally, because there are multiple GNSS constellations operating independently from one another, users can also combine the readings from multiple constellations to improve the accuracy of the measured time. The result of these multiple independent constellations is the global availability of accurate timing and synchronization, which is an incredibly useful and necessary resource for operators of critical infrastructure [2].

Despite the advancements and capabilities of GNSS, these systems should not be wholly relied upon to provide accurate timing and synchronization to critical infrastructure, for several reasons. First, because the satellites in these constellations are in medium Earth orbit (MEO), the signals they broadcast are weak by the time they are received at Earth’s surface. Weak signals are easier to jam and are easier to spoof, and therefore GNSS signals are subject to both problems, with no clear remedy. Second, because these satellites are in space, they are also subject to cosmic phenomena such as solar flares and impact from space debris, both of which have the potential to cause catastrophic damage. Finally, despite the robust history of accuracy and reliability of many GNSS, they are still subject to generic malfunctions, such as the GPS malfunction in January of 2016 [1] [3]. Malfunctions such as this can silently degrade the accuracy of the timing signals that are broadcast until the malfunction is rectified.

U.S. DEPARTMENT OF ENERGY | OFFICE OF ELECTRICITY

These problems, which can affect all GNSS constellations, demonstrate that without a reliable backup, over reliance on GNSS can become a single point of failure. This single point of failure can cause deleterious downstream effects to the various critical infrastructure that relies on stable timing and synchronization. Therefore, to avoid relying solely on GNSS, alternate timing and synchronization methods should be employed.

One of the few existing alternate synchronization methods that is accurate enough to serve as a GNSS backup and is also deployable to many geographically disparate users is known as two-way satellite time and frequency transfer (TWSTFT). In its simplest form, TWSTFT consists of two remote sites exchanging clock measurements with each other over a satellite communication link. At each remote site the received measurements are compared to the site's own measurements, which allows the time difference between the sites to be estimated [4]. This estimation can then be used to correct the local time offset relative to the remote site. As a result, so long as the remote sites are continuously exchanging data, both sites will remain synchronized to each other with a high degree of precision without utilizing GNSS. For example, the National Institute of Standards and Technology (NIST) has operated a TWSTFT link between their facility in Boulder, Colorado and the Physikalisch-Technische Bundesanstalt (PTB) in Germany since 2003, which allows NIST to contribute its time to the International Atomic Time (TAI) computed in France by the Bureau International des Poids et Mesures (BIPM) [5] [6]. However, the performance of a TWSTFT link often depends on many factors. Therefore, before recommending TWSTFT for widespread use as a GNSS backup or replacement, it is essential to create and thoroughly quantify and characterize the performance of a small, but geographically wide TWSTFT network over a substantial time period.

Experimental Design

To properly test the performance of a TWSTFT network, the network should consist of several individual sites that altogether form a wide geographical area. For this experiment, ORNL partnered with several public and private facilities to create the network. Specifically, facilities from Idaho National Lab (INL), National Institute of Standards and Technology (NIST), and Microchip Technology Inc. were utilized. The locations of these various sites, shown in Figure 1, demonstrate that the network covers a large portion of the continental United States and fulfills the desired geographical distance requirements for this experiment.

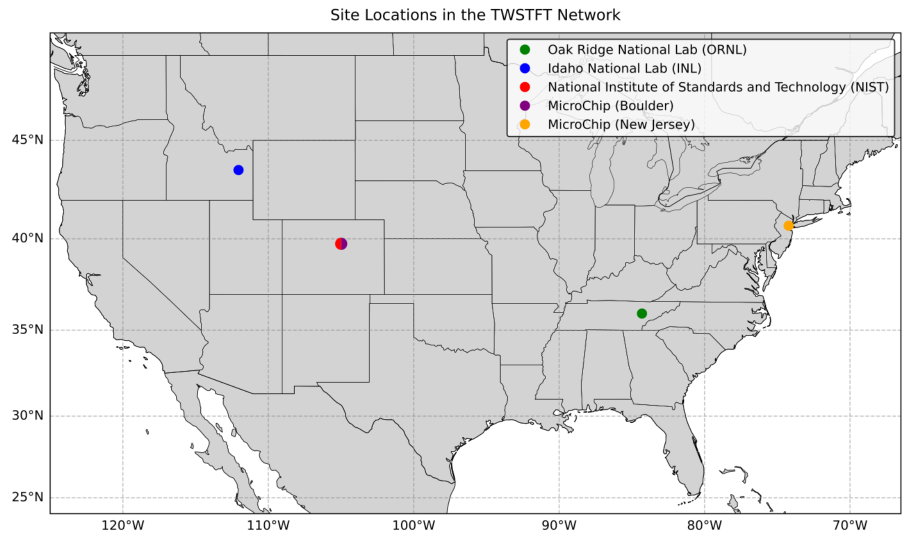

Figure 1. Locations in TWSTFT Network

Each site in the network contains several pieces of equipment to enable full participation in the network[1], such as a satellite ground terminal, a timing transfer modem, a reliable atomic clock (such as rubidium or cesium), and additional primary and backup timing sources. One of the most important pieces of equipment used in this implementation is the Microchip ATS 6502 Time Transfer Modem. This modem ingests a stable 10 MHz reference and a 1 Pulse-Per-Second (PPS) signal and then transmits these timing signals to the satellite ground terminal. Additionally, it also receives the same signals transmitted from other sites in the network and implements the differencing procedure described in [7] and outputs the difference in seconds between its clock and a remote clock. Each ATS 6502 modem can communicate with up to four other ATS 6502 modems and can also discipline its own time using the received time from one of these four remote clocks.

The goal of this experiment is to test the performance of time transfer over a TWSTFT network, and in addition test how well remote clocks can discipline each other over this network. While many modems will exchange timing signals, only the modems at INL and ORNL will discipline their respective clocks with each other's exchanged signals. Specifically, the modem located at INL will exchange timing signals with modems located at ORNL, NIST, Microchip Technology Inc. (New Jersey), and Microchip Technology Inc. (Colorado) and will discipline its clocks with the received timing signals from ORNL, while the modem located at ORNL will exchange timing signals with modems at INL, NIST, Microchip Technology Inc. (New Jersey), and Microchip Technology Inc. (Colorado) and will discipline its clocks with the received timing signals from INL.

After installing and configuring the required equipment at each site, the network was operational and entered a testing phase, during which its stability was monitored, and any software bugs/installation problems were corrected before entering the final experimental phase. Additionally, as described in [8], an additional piece of equipment required for a complete

[1] More in-depth information about the other necessary pieces of equipment for a complete TWSTFT installation as well as satellite details can be found in [6].

TWSTFT installation requires a GNSS connection to initialize the TWSTFT connection. Once the network was fully operational and stable, and immediately prior to the final experimental phase, the GNSS connection to this piece of equipment was severed at the ORNL and INL sites. This intentional disconnection was performed to reliably guarantee that the experimental results were not artificially impacted by a continuous GNSS connection.

Experimental Results

After severing the GNSS connection, data from each modem at ORNL and INL was collected for approximately ten days. The collected data was in the form of a time difference (in seconds) between that modem's clock and the clock of a remote modem, which was computed by the receiving modem. After collecting the bulk data from each modem, the individual data for each remote modem's received data was separated and processed. The data processing procedure for each bin of data was as follows. First, the median time difference and standard deviation over the entire collection period were computed. Next, using these two values, outlier points laying outside of the range

$$[-5|\sigma| + \mu, 5|\sigma| + \mu]$$

(where σ is the standard deviation and μ is the median value) were removed from the dataset. The point of this filter is to discard only extremely anomalous measurements from the plots but to keep almost all other measurements.

Next, remaining series of points were resampled using a 60 second average, that is, across the entire series, all points inside each 60 second window were averaged into a single point. This operation was performed to reduce the total data volume and make the resulting plots easier to interpret. Next, the median time difference computed earlier was subtracted from the remaining points in the set. This subtraction was performed to more clearly outline the total accumulated offset swing throughout the entire collection period. Finally, after removing the median difference, the finalized dataset was plotted, where the time differences were plotted on the y-axis and the dates were plotted on the x-axis as the fractional number of days since the start of the visualized data. For convenience, the start date and time is shown on the plot as a Modified Julian Date (MJD) and as well as an ISO 8601 timestamp. Plots showing the data recorded from the modem at ORNL and the timing signal from modems at INL, NIST, Microchip (New Jersey), and Microchip (Colorado) are shown in Figure 2, Figure 3, Figure 4, and Figure 5, respectively.

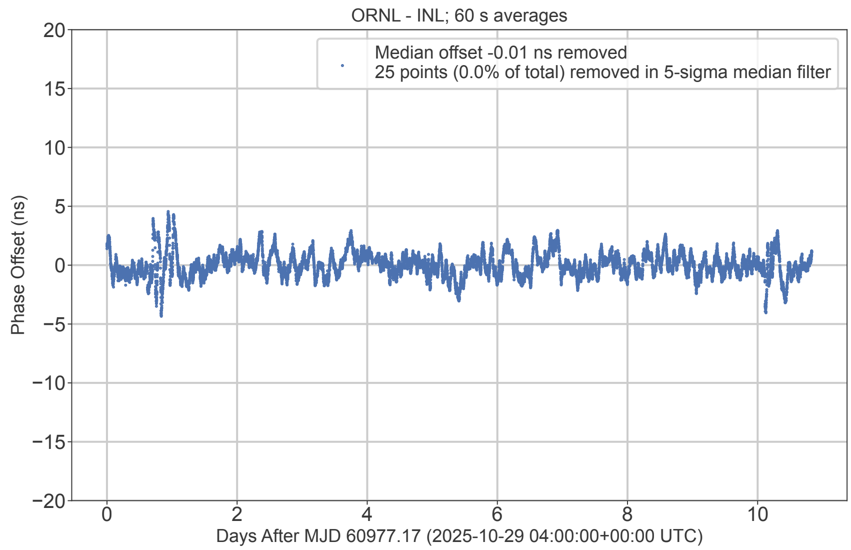

Figure 2. Time Difference Between ORNL and INL, Measured at ORNL

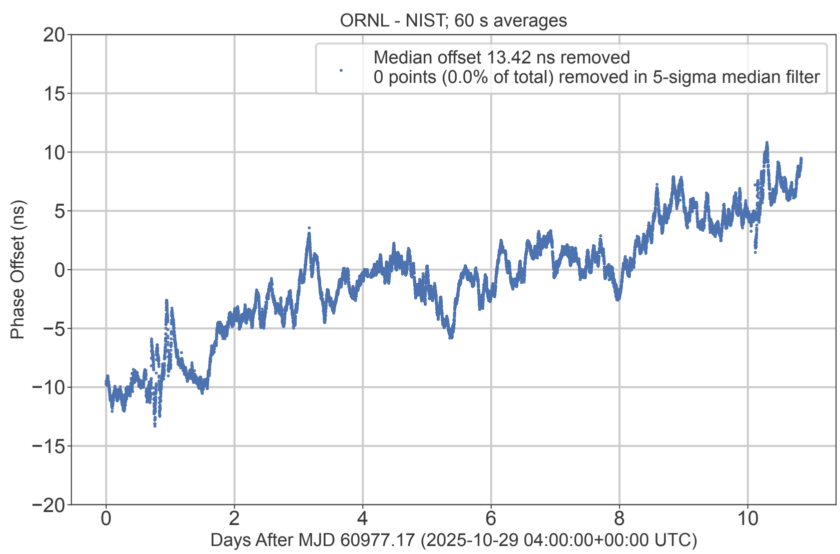

Figure 3. Time Difference Between ORNL and NIST, Measured at ORNL

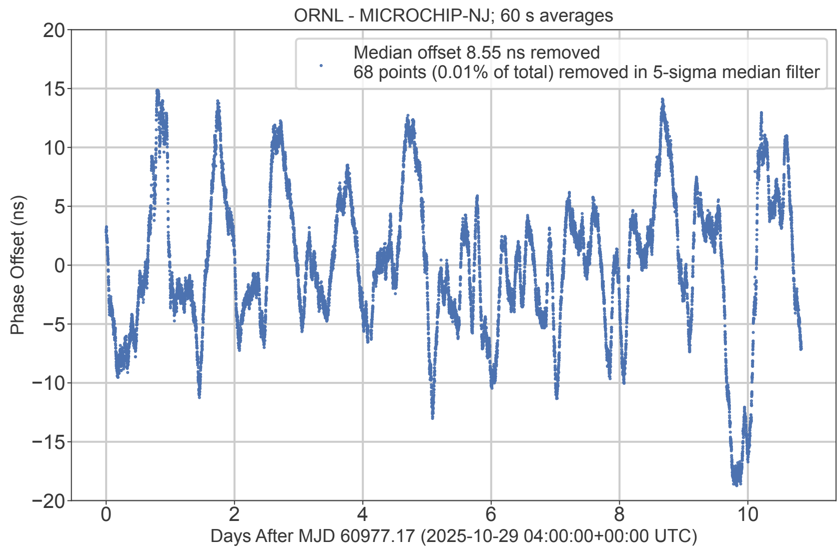

Figure 4. Time Difference Between ORNL and Microchip (New Jersey), Measured at ORNL

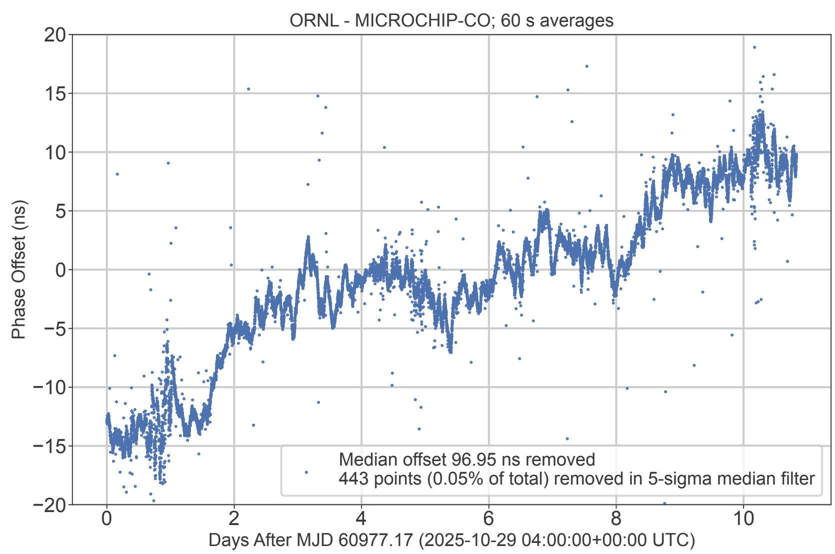

Figure 5. Time Difference Between ORNL and Microchip (Colorado), Measured at ORNL

Additionally, plots showing the data recorded from the modem at INL and the timing signal from modems at ORNL, NIST, Microchip (New Jersey), and Microchip (Colorado) are shown in Figure 6, Figure 7, Figure 8, and Figure 9, respectively.

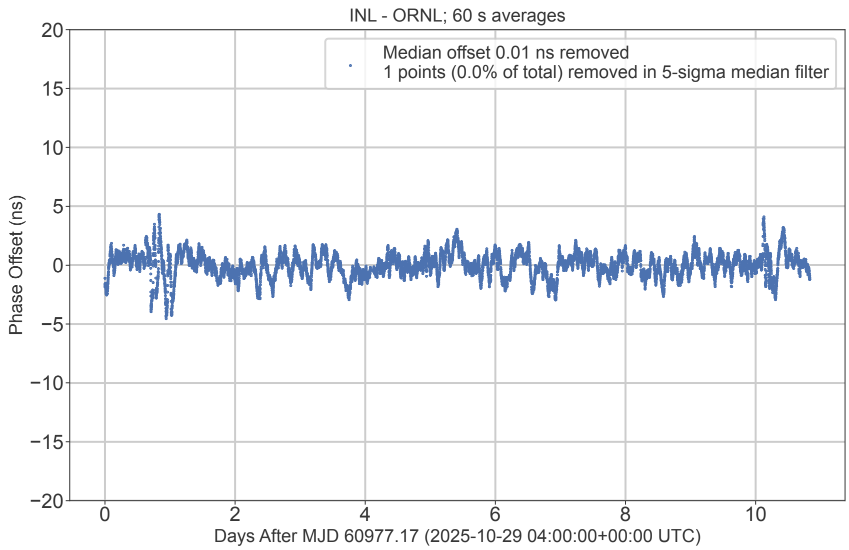

Figure 6. Time Difference Between INL and ORNL, Measured at INL

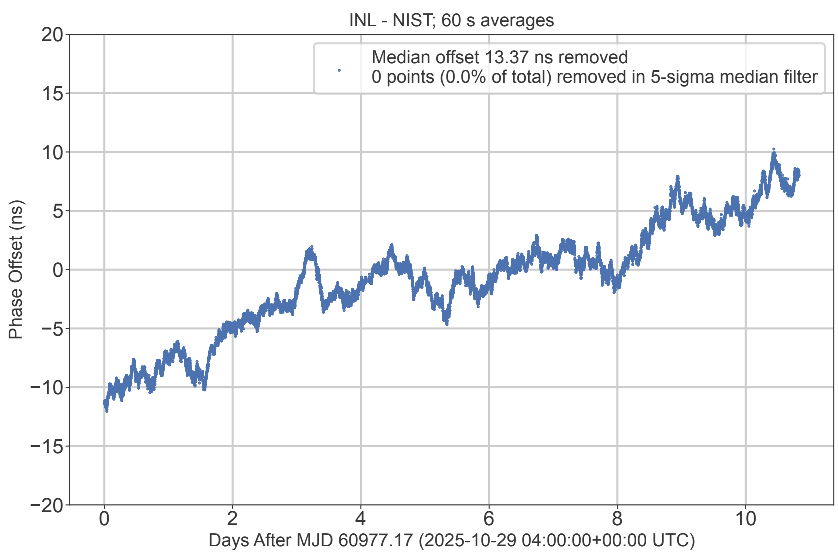

Figure 7. Time Difference Between INL and NIST, Measured at INL

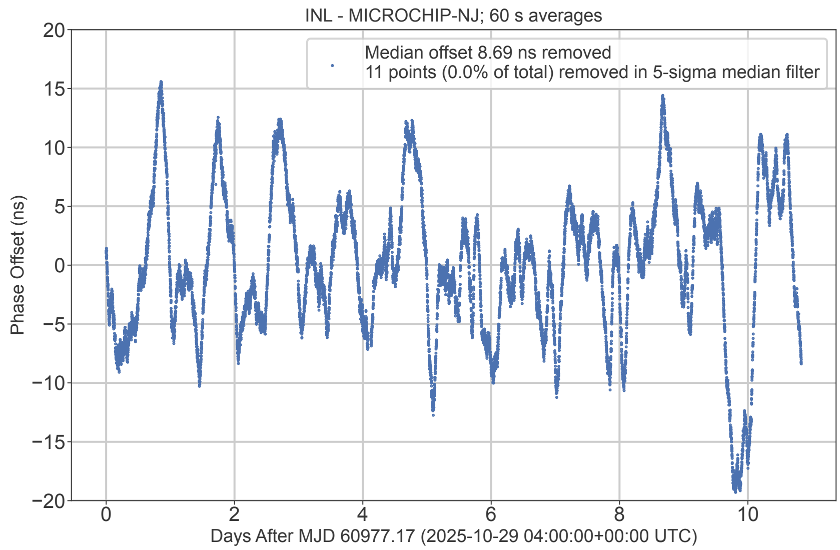

Figure 8. Time Difference Between INL and Microchip (New Jersey), Measured at INL

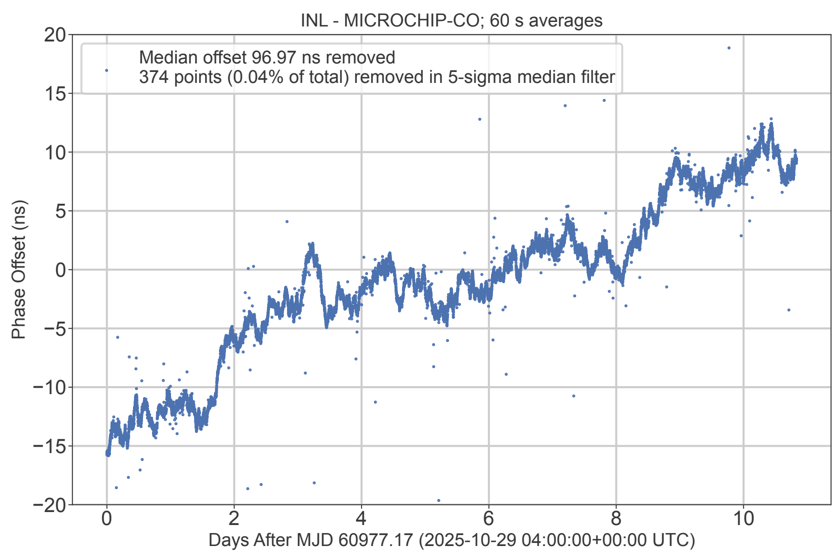

Figure 9. Time Difference Between INL and Microchip (Colorado), Measured at INL

Experimental Analysis

When first exploring the results, it should be noted that individual results from the same receiving location can be vastly different. For example, the ORNL-INL plot shown in Figure 2 demonstrates an offset with low variance that is relatively flat across the entire plot, while the ORNL-MICROCHIP-NJ plot shown in Figure 4 demonstrates a much more variable offset that oscillates between clear peaks and troughs. The shape difference between these and other results is caused mainly by slight hardware differences between locations and should not be interpreted as a sign of poor synchronization between sites.

For a more meaningful analysis of the results, the ORNL and INL data is examined first. As shown in Figure 2 and Figure 6, the time difference between the clocks at these two sites over ten days is close to zero with median differences of -0.01 ns and 0.01 ns respectively. This initial examination shows that the clocks at ORNL and INL remain very closely synchronized to each other (with +- 5 ns) over many days, which implies that each side is using received clock measurements to discipline its own clock. To further demonstrate this, when comparing the data in Figure 2 to the data in Figure 6, upon closer examination it is clear that each plot is a mirror of the other, that is, at any given point in time, the time difference is equal in magnitude but opposite in sign. This phenomenon can be shown more clearly in Figure 10, which is a plot of the ORNL and INL data plotted together.

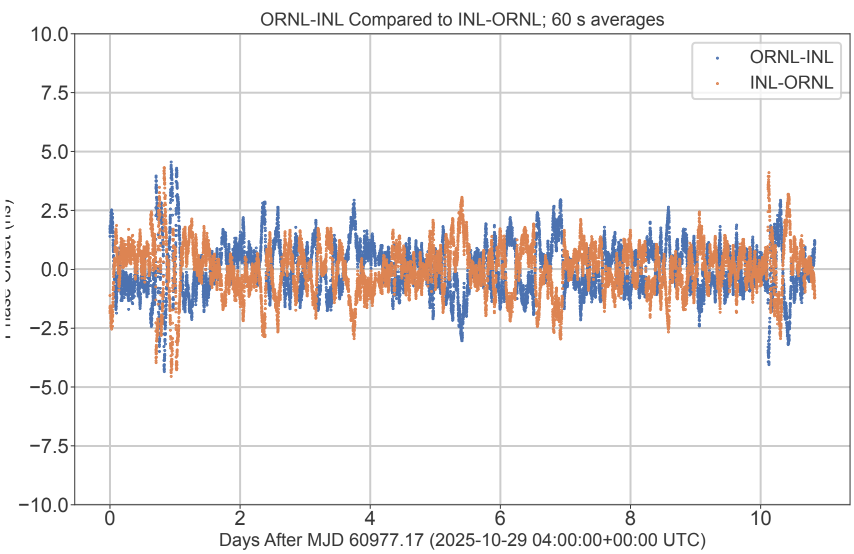

Figure 10. ORNL-INL and INL-ORNL Data Plotted Together

This is entirely expected (and desired) for a functioning time transfer network. Because both sides are measuring the difference in their time to the time from the other side, if time is being transferred successfully and both sides have nearly perfect knowledge of the actual clock measurements from the other side, then the time difference computed at one side should always be equal and opposite in sign to the difference computed at the other side. Additionally, this principle also applies to the median value, which aligns with the computed median values for each side shown in Figure 2 and Figure 6.

While the equal and opposite time difference plots partially validate that time is being transferred correctly, additional verification should be used as well. To show not only that time measurements are being correctly transferred across the network but also that two participants can use these measurements to discipline their own time, consider the separate plots comparing ORNL and INL to NIST in Figure 3 and Figure 7, respectively. If time is being correctly transferred and moved through the network, and if the clocks at ORNL and INL are correctly using this time to discipline their own clocks, then ORNL and INL should be nearly exactly synchronized to each other. If this is true, then the measurements comparing ORNL and INL clocks to the clock at NIST (which is not being disciplined by the network) should be nearly equal. When the two plots, shown in Figure 3 and Figure 7, are compared it is clear that the plots are nearly equal. This effect is more clearly shown in Figure 11, which plots the ORNL-NIST and INL-NIST data together. The plots are not only almost the same shape, but their median values are also nearly identical (13.42 ns measured at ORNL and 13.37 ns measured at INL).

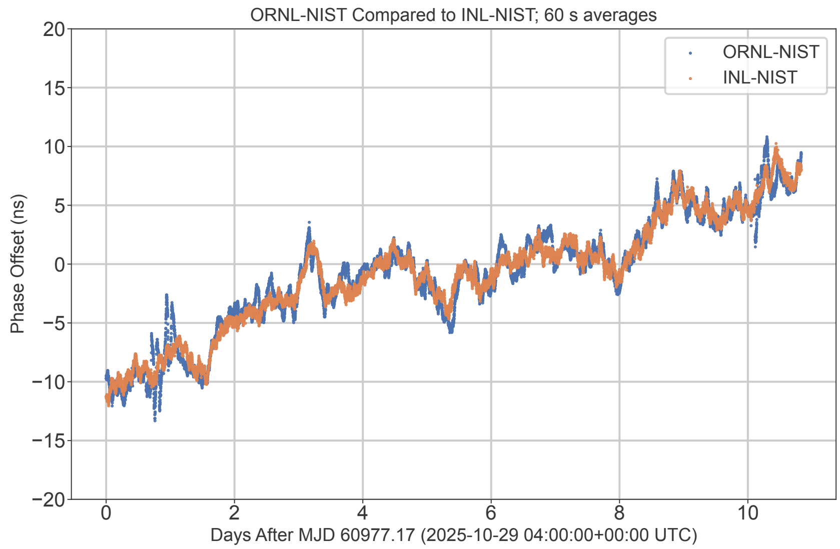

Figure 11. ORNL-NIST and INL-NIST Data Plotted Together

Repeating this same procedure for the modems at the two Microchip sites at Colorado and New Jersey produces the plots shown in Figure 12 and Figure 13. These plots both demonstrate the same behavior shown by the ORNL-NIST and INL-NIST plots, which is that the clocks at ORNL and INL are closely synchronized because comparisons to the same remote clock results in nearly equivalent time differences.

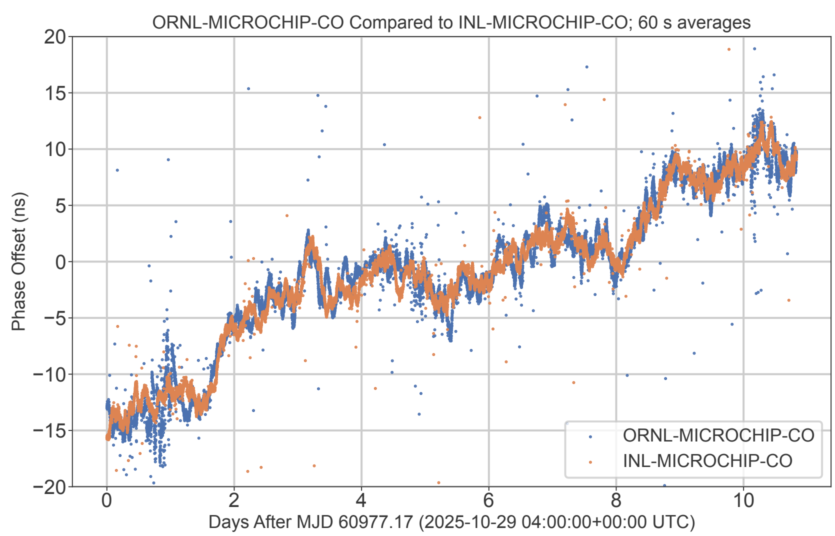

Figure 12. ONRL-MICROCHIP-CO and INL-MICROCHIP-CO Plotted Together

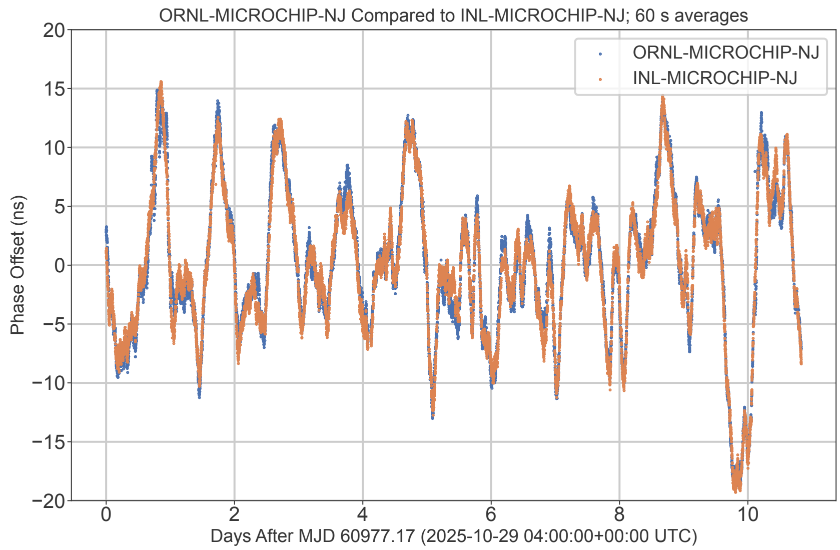

Figure 13. ORNL-MICROCHIP-NJ and INL-MICROCHIP-NJ Plotted Together

Conclusion and Future Work

In this experiment, a time transfer network was established over satellite, linking five different sites together. In this network, time was successfully transferred between several sites and the accuracy of synchronizing the clocks at ORNL and INL together using received time was quantified. When synchronizing these two sites, the offset at any given time was no more than five nanoseconds in magnitude over a period of approximately ten days. Additionally, when comparing the time difference individually between ORNL/INL and the clock at NIST and the clocks at the two Microchip sites, the plots were nearly the same, further demonstrating that the clocks at ORNL and INL are closely synchronized together. These results demonstrate that a TWSTFT network can synchronize clocks at remote sites to an accuracy comparable to that of a GNSS.

While these results demonstrate that clocks at different sites can be symphonized over satellite with an incredibly low discrepancy over a long time, there are several considerable caveats associated with the TWSTFT network built for this experiment. Firstly, the required equipment is expensive. The modem alone costs more than $50,000 and the network requires additional equipment to be fully operational. This makes the overall cost of implementing a TWSTFT network possibly cost prohibitive when many sites must be synchronized. This network also requires satellite bandwidth to be leased, which adds an additional and considerable recurring cost. While the individual cost per site for leased satellite bandwidth decreases linearly as sites are added to the network, it remains a major cost that must be shouldered. These combined costs make widespread adoption of TWSTFT networks unlikely, and future work should focus on any possible methods to reduce operating costs for equipment and/or satellite bandwidth. Finally, the ATS 6502 modem requires an initial GNSS connection to successfully initialize, even though it can operate without a connection after initialization. While this is a trivial problem when GNSS are abundant and fully functional, any phenomena that degrades or destroys the functional capacity of GNSS may also prevent or limit the function of a TWSTFT network unless the network has already been established and initialized. Future work that meaningfully addresses this severe limitation would dramatically improve the feasibility of TWSTFT networks to serve as a backup or replacement for GNSS.

Acknowledgements

We would like to acknowledge and thank both NIST and Microchip Technology Inc. for maintaining the TWSTFT setups at their respective facilities, which allowed us to compare our clocks at ORNL and INL to the clocks at NIST and Microchip.

References

[1] J. A. Sherman, "Time transfer alternatives to GNSS," Time and Frequency Division, NIST, 2024.

[2] U. S. Spaceforce, "Civilian applications of GPS – Timing," 2024. [Online]. Available: https://www.losangeles.spaceforce.mil/About-Us/Fact-Sheets/Display/Article/734551/. [Accessed 1 October 2024].

[3] J. Yao, M. A. Lombardi, A. N. Novick, B. Patla, J. A. Sherman and V. Zhang, "The effects of the January 2016 UTC offset anomaly on GPS-controlled clocks monitored at NIST," Time and Frequency Division, NIST, 2016.

[4] P. Banerjee and D. Matsakis, An Introduction to Modern Timekeeping and Time Transfer, Cham: Springer Nature Switzerland AG, 2023.

[5] V. Zhang and M. A. Lombardi, 4 March 2009. [Online]. Available: https://tf.nist.gov/general/pdf/2337.pdf. [Accessed 6 May 2026].

[6] Bureau International des Poids et Mesures (BIPM), 10 August 2021. [Online]. Available: https://www.bipm.org/documents/20126/59466374/6_establishment_TAR20.pdf/5b18b648-0d5a-ee02-643d-a60ed6c148fc. [Accessed 6 May 2026].

[7] D. W. Hanson, "Fundamentals of two-way time transfers by satellite," in *Proceedings of the 43rd Annual Symposium on Frequency Control*, 1989.

[8] UT-Battelle, LLC, "Lessons Learned for TWSTFT Implementations," UT-Battelle, LLC, Oak Ridge, TN, 2025.

The Center for Alternative Synchronization and Timing (CAST) at Oak Ridge National Laboratory (ORNL) performs research, development, testing, evaluation, and technical assistance to enabled resilient timing and synchronization for the power grid. Working closely with power utilities, timing hardware and software vendors, network operators, and federal stakeholders, CAST helps develop and validate alternative timing architectures to augment GPS time. CAST also translates and transfers ORNL's research and development (R&D) advances in secure timing and grid communications to power sector applications, and engages across the broader timing community to develop best practices to ensure the resilience of US critical infrastructure. CAST is sponsored by DOE's Office of Electricity. Visit https://cast.ornl.gov for more information.