Technical Bulletin #007: Holdover Experiment #1

CAST — Center for Alternative — Synchronization and Timing

ORNL/SPR-2025/3684

Summary: Stable and accurate timing forms the foundation of many modern essential services. While Global Navigation Satellite Systems (GNSS) provide a cheap way to achieve extremely accurate timing, these systems are vulnerable to interference and complete outages, forcing systems that rely on them into a state of holdover. Testing and characterizing holdover performance is therefore essential for ensuring that these modern services do not go down when GNSS-based timing is unavailable.

Introduction

This is the first in a series of three bulletins, each of which will demonstrate and explain the performance of various combinations of timing references while in a state of holdover. As shown in [1], this first experiment will test the holdover performance of a combination of a double oven crystal oscillator (DOCXO) and a magnetic cesium reference.

Experimental Design

For this experiment the Oscilloquartz OSA 5422, a compact grandmaster clock, will be used as the main clock being steered by the DOCXO and magnetic cesium references [2]. The specific magnetic cesium reference used for this experiment is the OSA 3230B. The OSA 3230B is a commercially available primary reference clock that uses magnetic cesium as its timing source. Figure 1 shows the complete experimental setup used for the duration this experiment.

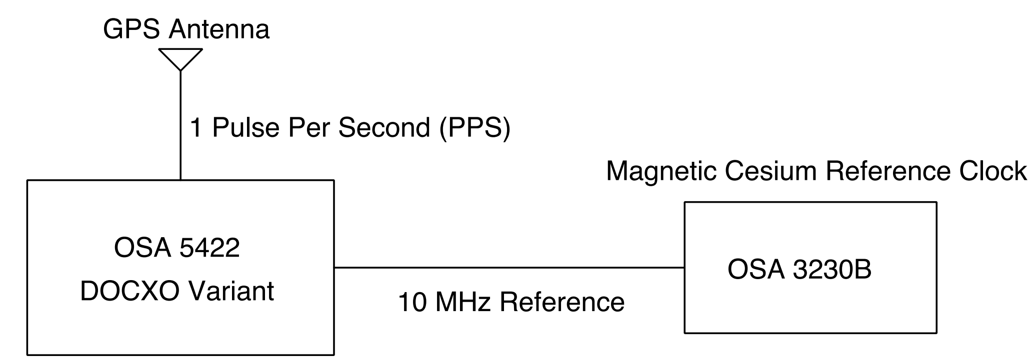

Figure 1. Experimental Setup Used in All Tests

To begin thoroughly characterizing the holdover performance of the OSA 5422 clock being steered by the previously specified references, a total of three tests were performed. Prior to each test, the OSA 5422 was first allowed to achieve a GPS lock and synchronize its time. To

U.S. DEPARTMENT OF ENERGY | OFFICE OF ELECTRICITY

allow the OSA 5422 sufficient time to settle into a steady state, the OSA 5422 was left to be steered by GPS, the onboard DOCXO, and the OSA 3230B magnetic cesium reference for a period of roughly 30 days. Once this period had concluded, GPS steering was disabled. At this stage, the clock officially entered holdover, and its drift, also known as the phase offset from GPS time, was measured and recorded once every second for the duration of the test. For the three tests performed, the first test was conducted over a duration of 14 days, the second test was conducted over a duration of 30 days, and the final test was conducted over a duration of 60 days. It should be noted that these three tests were discrete and conducted separately from each other.

Test Results

Following the conclusion of the final test, the results of each test were collected and processed. For each test, the phase offset measurements were decimated by a factor of 300 by averaging every five minutes. Five minutes was chosen as the averaging window because the clock only creates and transfers a single file containing the one-second measurements every five minutes, and thus a five-minute averaging window was easy to implement. In addition, there were gaps occasionally present in the recorded data due to file transfer failures. When these gaps occurred, the results were filled in using linear interpolation. To make these gaps apparent and to not further errant conclusions from interpolated data, all gaps greater than one hour have been annotated using shaded regions on each plot.

The first test, which lasted a total of 14 days, ran from January 25, 2024 until February 8, 2024 and is shown in Figure 2.

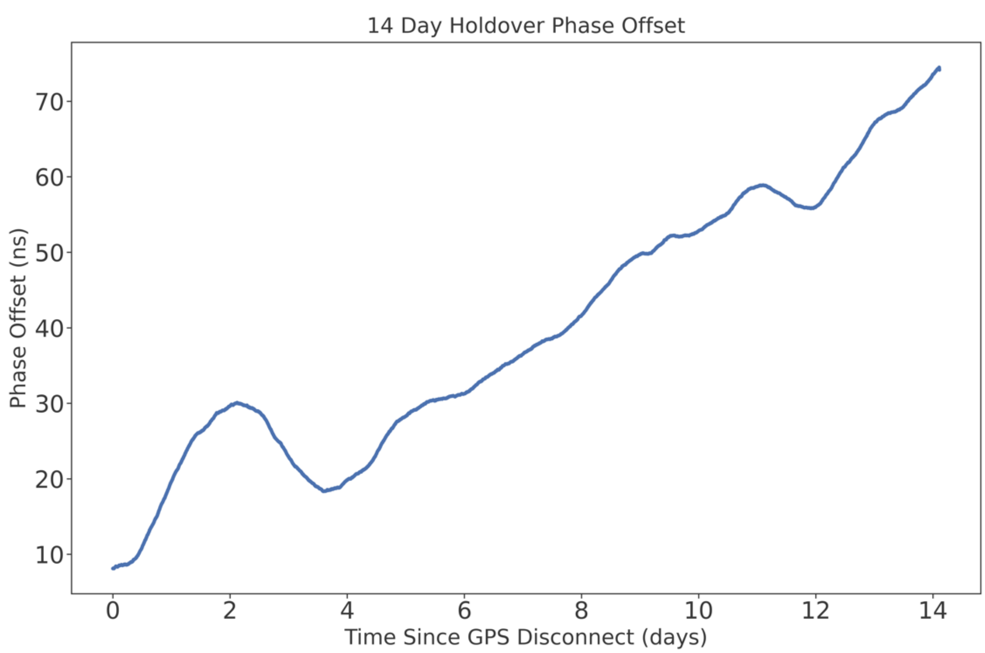

Figure 2. Phase Offset Results from 14 Day Holdover Test

The second test, which lasted a total of 30 days, ran from March 18, 2024 until April 17, 2024 and is shown in Figure 3.

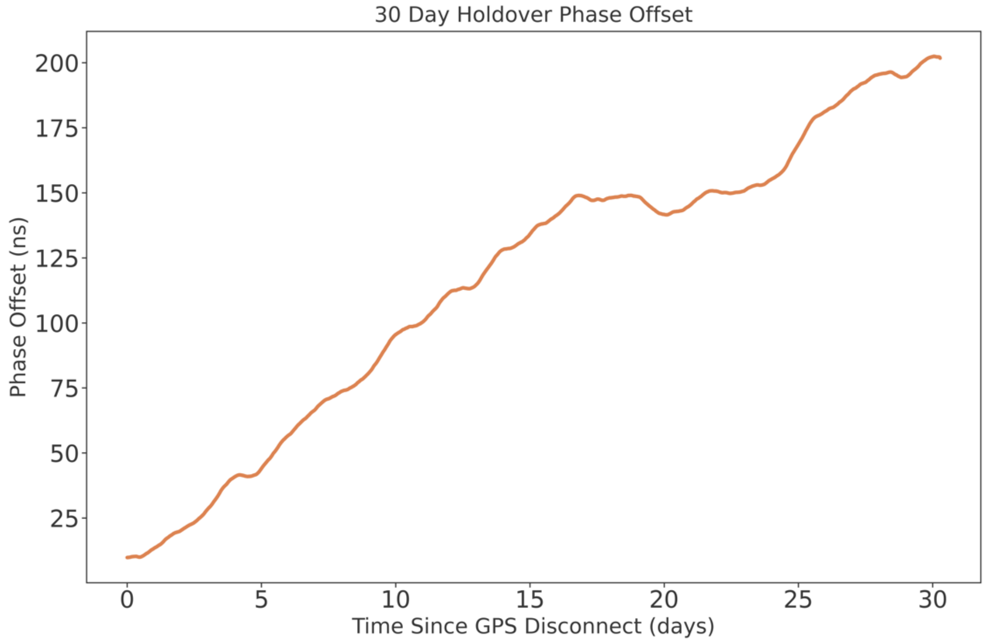

Figure 3. Phase Offset Results from 30 Day Holdover Test

The third and final test, which lasted a total of 60 days, ran from June 18, 2024 until August 17, 2024 and is shown in Figure 4.

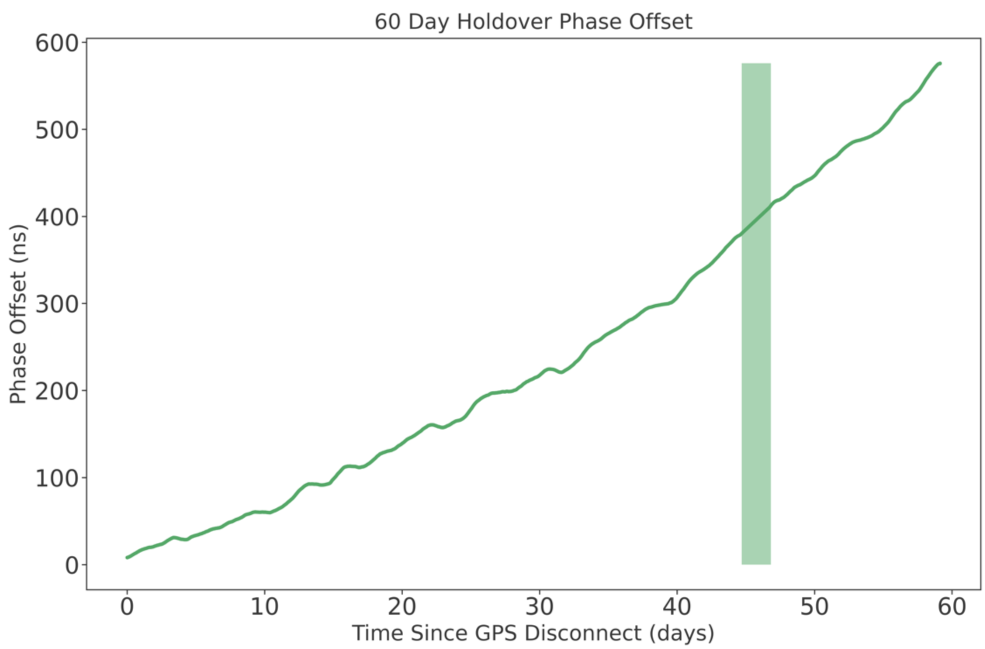

Figure 4. Phase Offset Results from 60 Day Holdover Test

Unlike the first and second tests, this third test included a large gap, as shown in the shaded region between the 40- and 50-day mark. The gap shown in Figure 4 lasted approximately 2 days in total duration. Despite this gap, Figure 5 demonstrates that when the results from each test are plotted together, all three trend loosely in the same direction.

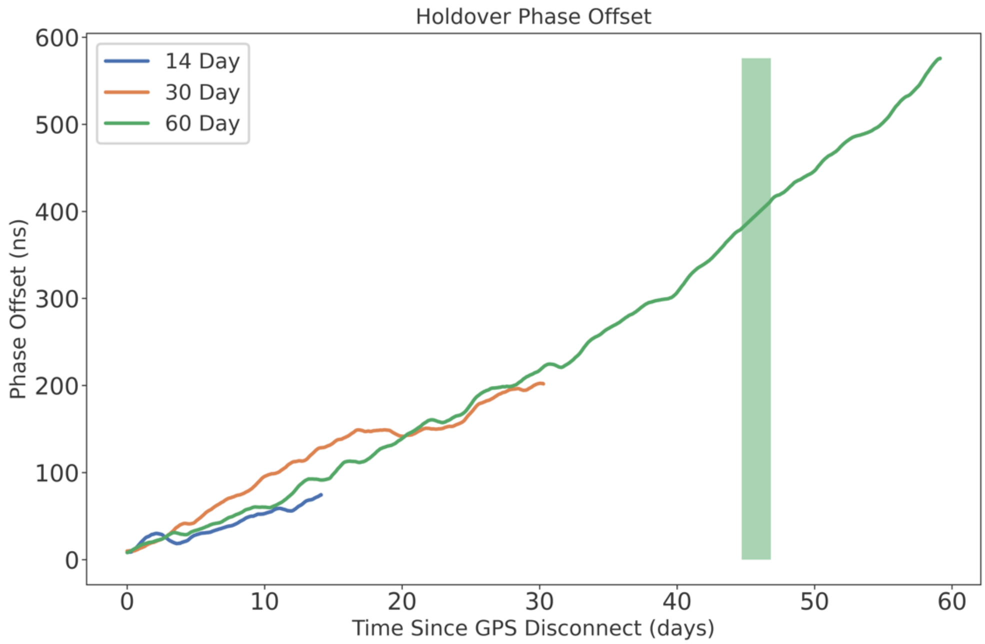

Figure 5. Phase Offset of all Three Holdover Tests

While it is trivial to discern via inspection that each test is trending in the same direction relative to the others, it is more difficult to discern how the mean phase offset change in nanoseconds per day of each test compares to the other tests. Characterizing the average phase offset change in nanoseconds per day over a long testing period is a useful way of distilling a large amount of data into a single metric. This single metric can then be used to compare the holdover stability of one clock to the holdover stability of another. Table 1 depicts the average phase offset increases in nanoseconds per day for each of the holdover tests. In addition, each test was also divided into equal periods of 14 or 15 days[1], and the average phase offset increase for each of these equal periods was calculated and is enumerated in Table 1.

Table 1. Average Phase Offset Increase in Nanoseconds per Day

Test	Total Average Increase (ns/day)	First Period (ns/day)	Second Period (ns/day)	Third Period (ns/day)	Fourth Period (ns/day)
14 Day Holdover	4.73	4.73	-	-	-
30 Day Holdover	6.34	8.25	4.56	-	-
60 Day Holdover	9.40	5.94	7.97	11.07	13.35

Overall, the average phase offset change appears to increase as the total holdover duration increases when analyzing the results in Table 1. This result implies that the change in phase

[1] Only the first test ("14 Day Holdover") refers to a period of 14 days. The other two tests consist of average increase calculations over 15-day periods.

offset is non-linear. To explore this possibility more thoroughly, a combined rate of change plot was created and analyzed. To create this plot, first the rate of change with respect to time in days was computed over each plot shown in Figure 5. The resulting rate of change plot, shown in Figure 6, is difficult to analyze due to the presence of noise. To filter out the noise and smooth the rate of change plot, a windowed moving average with a variable length was applied to the plot. The length of the window was set to approximately 25 percent of the total number of samples in each test. This percentage was determined experimentally to balance a tradeoff between the smoothing effect of the moving average and the loss of features as the window size increases. Setting the window size to 25 percent of the total number of samples eliminated most of the undesirable noise but preserved important features.

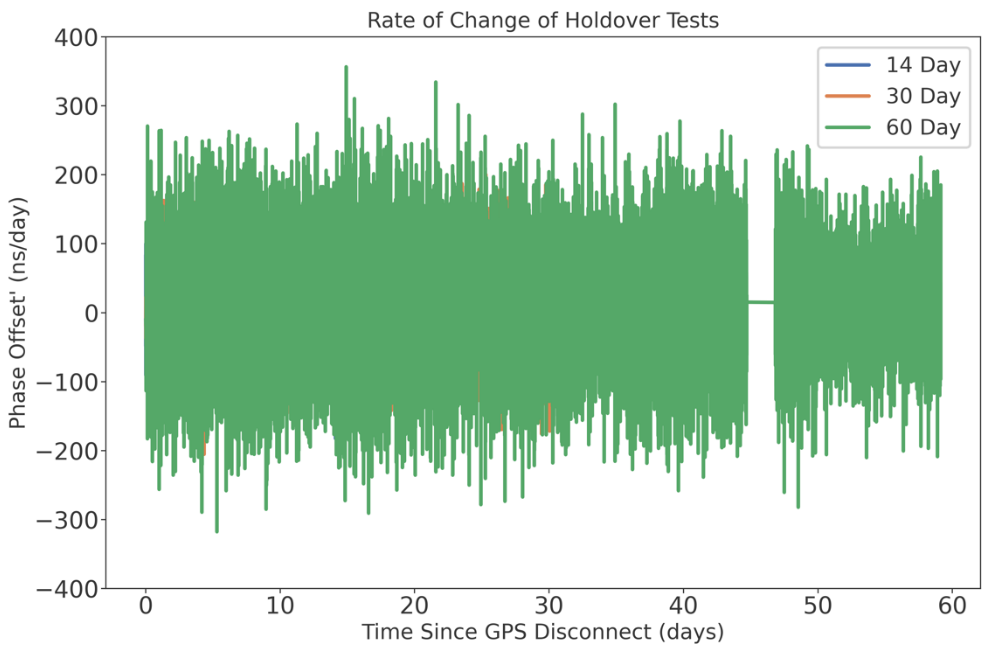

Figure 6. Phase Offset Rate of Change for All Holdover Tests

Shown in Figure 7, the rate of change plot with a windowed moving average is cleaner[2] than the previous plot and confirms that the phase offset change in nanoseconds per day increases as the total holdover duration increases. In summary, these results appear to demonstrate that the phase error of a clock in a state of holdover is non-linear. Instead, the average phase error increases as a clock spends more time in holdover. This finding may be significant because non-linear functions can increase more rapidly with respect to time compared to linear functions.

[2] It should be noted that as a moving window average slides across and begins to reach the end of the data, the window starts to "slide off" of the data, resulting in sudden changes at the end of the plot relative to the rest of the plot. This is caused by the point averages consisting of fewer total points as the window continues to slide off. To minimize this effect, any point averages that contained fewer than 50% of the total number of required points specified by the moving average window size were removed from the plot in Figure 7.

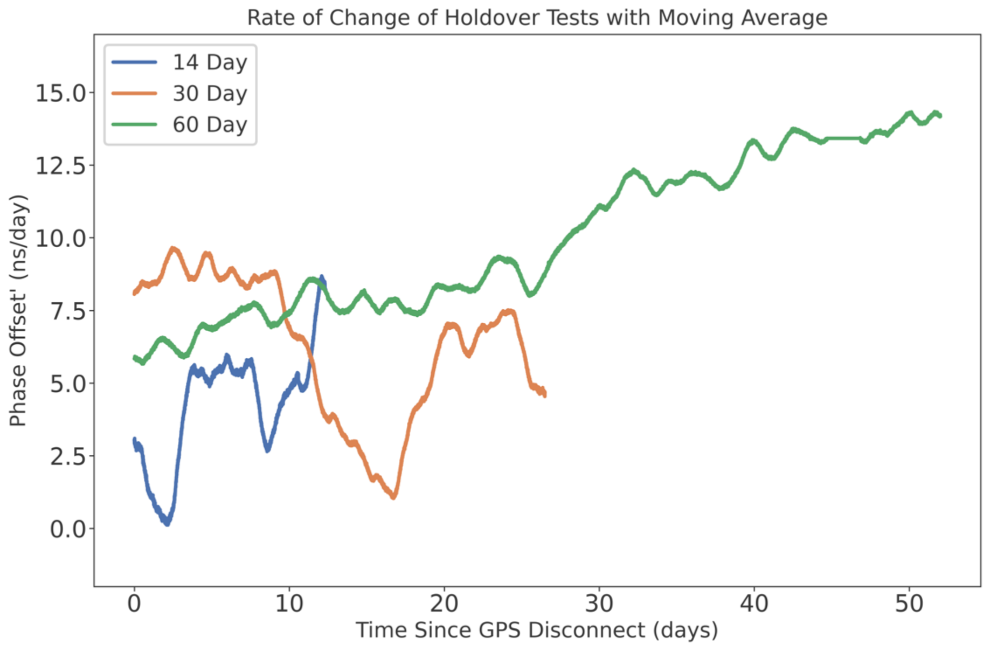

Figure 7. Phase Offset Rate of Change for All Holdover Tests with Moving Average

Conclusion and Future Work

In this set of experiments, the holdover performance of an OSA 5422 grandmaster clock steered by an onboard DOCXO and an OSA 3230B magnetic cesium reference source was tested and characterized. Overall, the holdover performance of this experimental setup was excellent, with average phase error increases between four and ten nanoseconds per day. While the holdover performance of this setup doesn't quite compare to the cesium fountain clocks operated by the NIST [3], it is more than enough to satisfy the timing requirements of terrestrial telecommunications standards such as the 5th generation (5G) mobile network standards defined by the 3rd Generation Partnership Project (3GPP). 5G, being the newest currently deployed wireless mobile network standard, defines strict timing requirements for some modes of operation. For example, the time division duplexing (TDD) mode of 5G requires that the timing of the base stations and handsets to be no more than 1.5 microseconds offset from each other [4]. Even if only ten percent of that margin were available, this remaining margin is still equivalent to 150 nanoseconds, or roughly a duration of 20 days in holdover using this experimental setup, which is more than enough to cover even very long GNSS outages.

One area that future experiments will address would be the operating cost of this clock setup. While this experimental setup is performant and easy to deploy, it may be cost-prohibitive to operate, as the complete stack as shown in Figure 1 costs roughly $70,000. The bulk of this cost is the magnetic cesium reference clock and because of this prohibitive cost, cheaper options will be explored in future work. To determine if cheaper options are viable, the holdover performance of experimental setups that utilize cheaper rubidium and optical cesium reference clocks will be characterized in the same manner during later experiments.

References

[1] UT-Battelle, LLC, "Holdover Experimental Summary," UT-Battelle, LLC, Oak Ridge, TN, 2024.

[2] Adtran, Inc., "OSA 5422 Compact PTP grandmaster, NTP server, SB/MB-GNSS receiver, multi-interfaces," 2024. [Online]. Available: https://www.oscilloquartz.com/-/media/oscilloquartz/resources/data-sheets/pdfs/osa-5422.pdf. [Accessed 03 December 2024].

[3] J. Sherman, "How UTC(NIST) Works," NIST, 20 November 2020. [Online]. Available: https://www.nist.gov/pml/time-and-frequency-division/time-realization/utcnist-time-scale-0/how-utcnist-works. [Accessed 2 October 2024].

[4] VIAVI Solutions Inc., "What is 5G Timing and Synchronization?," 2019. [Online]. Available: https://www.viavisolutions.com/en-us/what-5g-timing-and-synchronization. [Accessed 3 October 2024].

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.