GPS Solutions (2021) 25:141
https://doi.org/10.1007/s10291-021-01176-1

ORIGINAL ARTICLE

Benefits of operating multiple atomic frequency standards for GNSS satellites

John P. Janis[1] · Michael R. Jones[1] · Nicholas F. Quackenbush[1]

Received: 16 October 2020 / Accepted: 27 August 2021 / Published online: 9 September 2021

Abstract
For decades, Global Navigation Satellite System (GNSS) payloads have incorporated multiple Atomic Frequency Standards (AFS) to provide both redundancy and technology diversity. As space borne AFS have matured, they have demonstrated sufficient reliability to allow multiple AFS to remain active without significantly impacting mission life. Through the use of precise onboard measurements of the relative phases of the active AFS and a voltage controlled crystal oscillator an ensemble of the onboard clocks may be realized. For payloads with three active AFS an improvement in stability, and thus the clock component to User Range Error (URE), of a factor of 1.7 is realized for all averaging times beyond the phase locked loop time constant. In addition, it has been shown that while overall AFS reliability has been excellent, they have experienced occasional anomalous frequency and phase shifts throughout their life requiring control segment intervention. The operation of multiple AFS allows monitoring each of the clocks with respect to the other and autonomously correct for these anomalies. An evaluation of the key parameters governing the performance of these algorithms, along with their impact on URE, will be presented.

Keywords Timing · Stability · GNSS · Timekeeping · Allan deviation · Ensemble · Clock · Resiliency · GPS · Galileo · Glonas

Introduction

Since its inception, the fundamental architecture of GNSS has relied on the stability of an onboard AFS to provide global Positioning, Navigation, and Timing (PNT). This architecture choice minimizes the requirements for the user receiver but necessitates the inclusion of a high stability AFS. From the start, this architecture was plagued by the difficulty in designing AFS that operate completely unattended for years in Middle-Earth Orbit (MEO) with the associated environmental extremes. The early AFS used for GNSS experienced relatively high failure rates when compared to the desired mission life. To compensate, standard GNSS payload configurations have included multiple AFS, cross-strapped so that any single unit could be operated as the primary reference for the satellite.

✉ John P. Janis
John.Janis@L3Harris.com

1 Space and Airborne Systems, L3Harris Technologies, Clifton, NJ, USA

Modern GNSS AFS have now achieved extremely high levels of reliability as evidenced by the 20 years of failure-free operation achieved by a Rubidium Atomic Frequency Standard (RAFS) in 2017 on the first GPS-IIR vehicle, Taylor (2017). We will discuss the advantages of alternate architectures where multiple AFS are operated simultaneously and ensembled to provide the reference clock for the GNSS payload. While the marginal reliability benefits of multiple cold spares are lost in this scheme, the ability to meet all mission requirements should any or all of the additional AFS fail is maintained.

We will first describe the architecture of a generic Timekeeping System (TKS) necessary to leverage the benefits of multiple active AFS. This discussion will include some of the key design considerations and resultant performance. We then discuss the implementation of simple ensembling of the onboard AFS and the resultant performance benefits for the generic example of a system with 3 AFS. Finally, we present the advantages of operating multiple AFS in providing resiliency to clock anomalies generalizing the work presented in Janis and Jones (2017). As we will show, the benefits to the user are significant, providing both reduced URE and improved resilience in the face of clock anomalies.

TKS overview

While the output of the onboard AFS can be used directly as the system clock for a GNSS navigation payload the inclusion of a TKS provides significant flexibility in the design. The role of the TKS is to provide both control and monitoring of the system clock via both Ground commands and onboard Flight Software (FSW). This disciplined system clock is then used as the basis for all navigation signal generation. The functions and the algorithms used in the TKS are hosted in the FSW and may be reprogrammed on-orbit as new capabilities or unique Space Vehicle characteristics are discovered.

The basic architecture for a generic GNSS TKS is shown in Fig. 1. The TKS accepts several reference clock inputs, in this case three AFS, but any available stable frequency source could be substituted. The system clock is a controllable frequency source with excellent phase noise and short-term stability characteristics. In this example, an ultra-stable Voltage Controlled Crystal Oscillator (VCXO) nominally operating at the desired system clock frequency. All the clock sources are provided to a Precision Phase Measurement (PPM) system that generates digital phase measurements relative to a common, independent reference integrated into the PPM system and not shown in Fig. 1. The critical aspect of the PPM is to provide low latency, high accuracy and low noise measurements of the relative phase difference between each of the input clock sources. The sample rate for a generic digital phase measurement

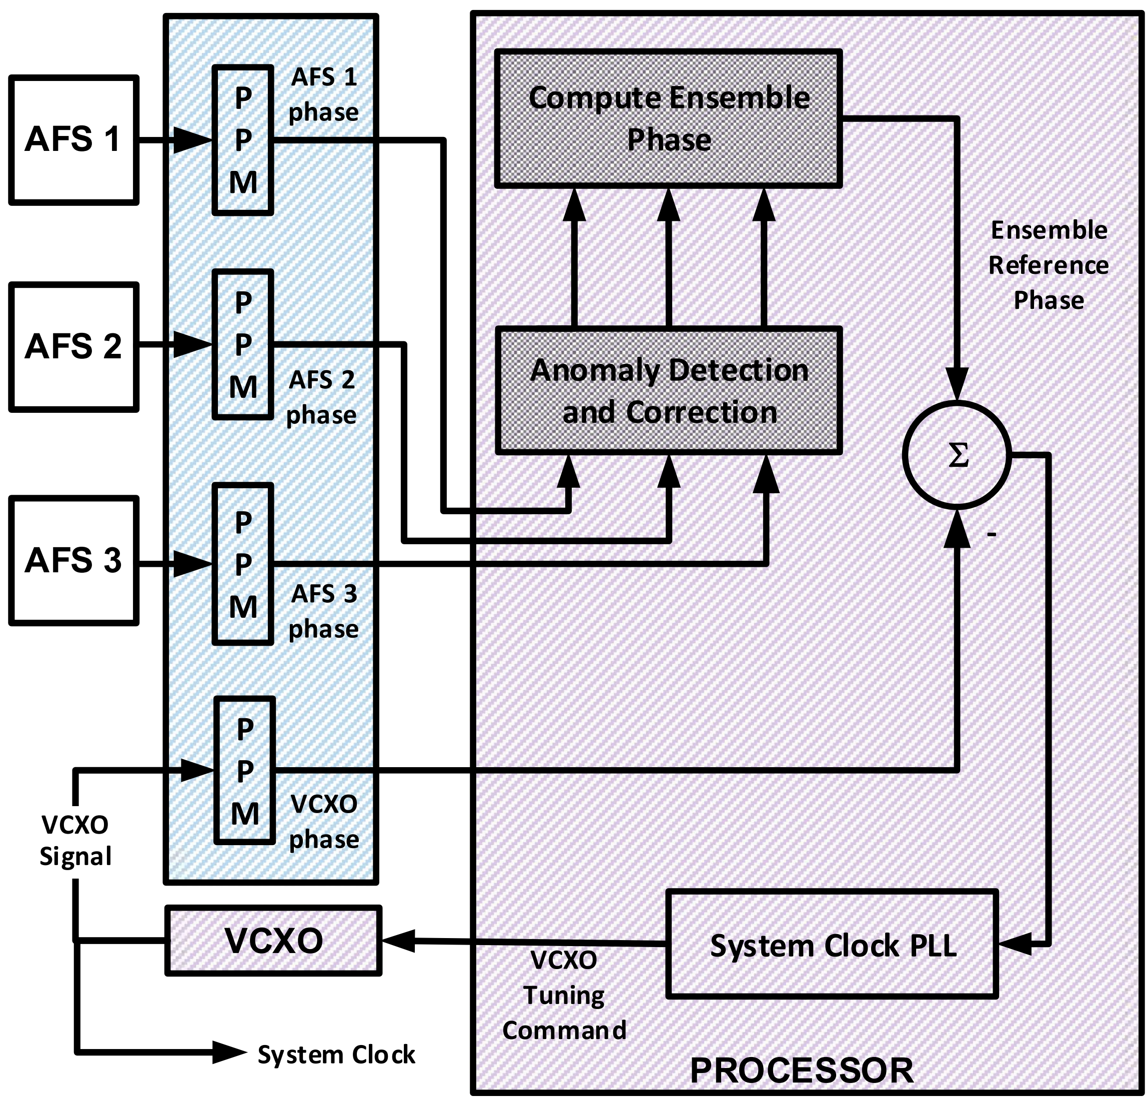

Fig. 1 Generic Timekeeping System Block Diagram: *A VCXO system clock is phase locked to a combination of AFS* via *digital phase measurements and PLL*

approach described in this paper is chosen for the best balance between reducing noise in the measurements to meet timekeeping system accuracy requirements, the bandwidth of the system clock Phase Locked Loop (PLL) that ingests the measurements and internal hardware delays.

The raw PPM phase measurements are made available to a flight computer for final processing to produce the system clock PLL error signal. The baseline configuration is to phase lock the high short-term stability VCXO to the high long-term stability AFS. This is implemented through a digital phase locked loop that minimizes the measured phase difference between the VCXO and AFS. VCXO frequency tuning commands are generated through a loop filter with bandwidth tuned to the crossover point between the VCXO and AFS stability. It is important for proper digital PPM stability to minimize the hardware delay between the time that the phase measurements are collected and the realization of the VCXO frequency change to maintain lock to the AFS.

For a generic GNSS TKS we will consider typical RAFS performance to lie in the range of the GPS IIR and GPS IIF performance as presented in Camparo et al. (2016) and a nominal high performance VCXO. Typical parameters for a generic system are shown in Table 1 with the AFS and VCXO stabilities characterized by their Allan Deviation at an averaging time of 1 s. Thus, this generic TKS provides access to relatively high rate and high precision phase measurements of each of the onboard clocks coupled with a software-controlled combination of those measurements capable of controlling the GNSS system clock.

For the baseline generic TKS performance, the phase measurements from a single AFS are passed directly through the Anomaly Detection and Correction (ADC) and the ensembling blocks. The phase difference between the VCXO and the chosen AFS are then fed into the digital PLL to generate VCXO tuning commands. The resulting simulated system stability, along with the input clock performance, can be found in Fig. 2. The system clock preserves the stability of the VCXO for short averaging times then takes on the stability of the AFS for long averaging times, with the crossover happening for averaging times equal to the loop time constant.

Table 1 Generic TKS Parameters

TKS Parameter	Value	Units
AFS White Frequency (@1 s)	2E-12	
VCXO Flicker (@1 s)	5E-13	
Phase Measurement Noise (White PM)	5E-12	s
Phase Measurement Interval	500	ms
Loop Time Constant	20	s
Loop Delay	1	s

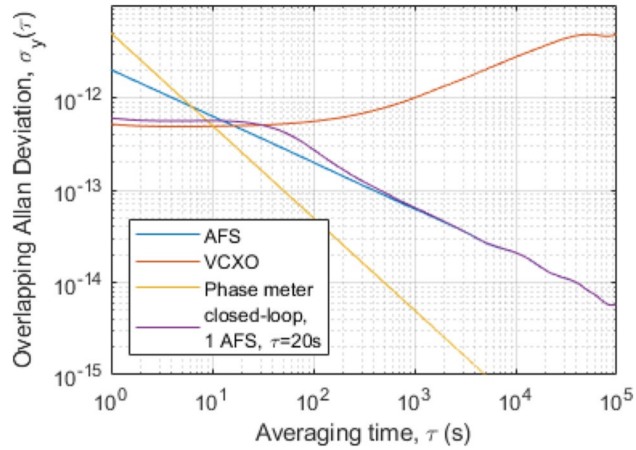

Fig. 2 Generic TKS Stability: *Closed loop TKS Allan Deviation performance when locked to a single AFS with a 20 s loop time constant*

Clock ensembling

In principle, any GNSS payload that operates multiple, independent reference clocks can employ ensembling to improve system clock stability. Here, independent is defined to mean the stochastic noise of the multiple reference clocks is statistically independent. Ensembling combines the signals of the multiple clocks through some averaging process to reduce the random noise component that contributes to stability (excluding environmental effects). In many GNSS applications, the goal is to improve the long-term stability of the system clock. Doing so alleviates the need for frequent ground updates to maintain the accuracy of the user navigation message, which contains a reference to the larger GNSS timescale maintained by the control segment. Ensembling can also be viewed as a means of enhancing GNSS robustness in as much as the system can be autonomous, i.e. with no ground contact, for a longer period of time while still providing acceptable user PNT performance.

Different approaches to ensembling can be used, ranging from simple weighted averaging of the measured phase at each epoch to a full Kalman filter implementation. For on-orbit applications, simpler is usually preferable to keep the burden on onboard computational assets as low as possible. The TKS architecture depicted in Fig. 1 is advantageous for simple, efficient ensembling and is realizable in practice. The three reference AFS are of the same technology, manufacturer and built to the same requirements. Hence, it is reasonable to assume all three have comparable stability and $1/\sqrt{\tau}$ white FM noise characteristics. Simple, weighted averaging is feasible with this architecture. We acknowledge that a system with a mix of clock types and technologies would require more sophisticated weighting. This weighting may depend on timescale if the noise statistics are different, e.g., combining $1/\sqrt{\tau}$ white FM with τ^0 flicker.

Ensembling is realized through a weighted average of the three independent AFS phase measurement streams from the precision phase meters in this notional implementation. The weights are fixed but can be configurable if desired, and are subject to the constraint that all weights must sum to 1. Implementation of ensembling in this architecture only requires the addition of software routines with no hardware modifications. The modifications to the FSW are added upstream of the system clock PLL and VCXO control loop, ensuring that system stability is not impacted.

Ensembling of the three reference AFS with a weighted average of their phase is given as

$$u_{en} = w_1 u_1 + w_2 u_2 + w_3 u_3 \quad (1)$$

where u_x is the phase of each AFS input as measured by the PPM against a common reference. For the case where the clocks are of the same type and performance $w_1 = w_2 = w_3 = 1/3$. However, a two-clock combination, for example, the first two of the three AFS, is trivial to implement by assigning $w_1 = w_2 = 1/2$ and $w_3 = 0$. A fall-back operation to just one clock of the three is achievable with a weight of 1 assigned to any clock of the three and 0 to the other two. All these modes of operation are options that require no changes to hardware or the need for special operational procedures.

An example of how simple, fixed weight ensembling can enhance system clock stability is shown in Fig. 3. A laboratory experiment was set up as shown in Fig. 1, with three independent clock inputs each with the same Allan Deviation performance as shown in Table 1. The blue curve is the measurement of system clock stability with just one AFS as the reference. This is equivalent to a special case of ensembling with one clock of the three given a weight of 1. The red curve is the measured system clock stability when locked to the ensemble of three AFS combined according to (1) with equal weighting. For averaging times shorter than the loop time constant, where the PLL is only loosely coupled to the AFS ensemble, the inherently superior stability of the VCXO dominates. Note that the PLL time constant is reduced for the ensembled case to take advantage of the earlier crossover point between the more stable ensemble and the VCXO. For averaging times longer than the PLL time constant, the system clock is tightly coupled to the AFS ensemble and its stability mirrors that of the ensemble, including the $1/\sqrt{\tau}$ noise statistics of the ensemble. The resulting improvement in stability over a very wide range of timescales is better than 40%, closely matching our theoretical model, and it is those longer timescales that matter for GNSS PNT robustness. For very long times of the order of days, the improvement may depend on the details of clock behavior in the presence of environmental effects and the possibility that, say, a residual random walk component of noise may occur. Nevertheless, this example illustrates that ensembling is indeed feasible for a PNT payload architecture that supports multiple references by design with only an FSW enhancement.

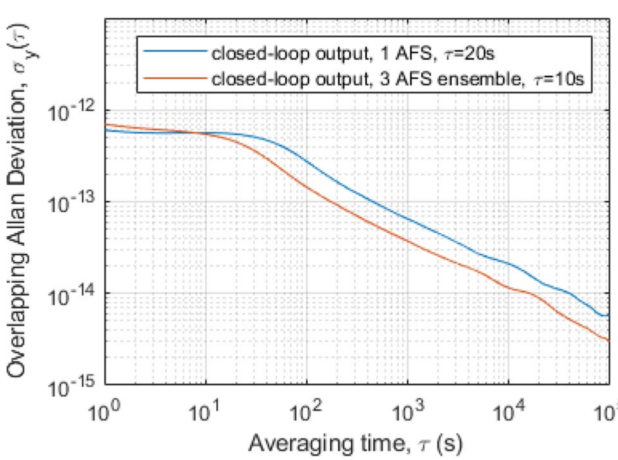

Fig. 3 TKS Stability with 3 AFS Ensembling

Anomaly detection

As in Formichella (2016), it has been shown that AFS on-orbit demonstrate a non-stochastic frequency jump phenomenon that is well described by a compound Poisson process. This process is the result of discrete but unpredictable jumps in frequency and is clearly distinguished from the more traditional continuous Wiener process behind the random walk. The discrete and random nature of these jumps allows for them to be distinguished from the background noise processes when compared to AFS that have not experienced a jump. We present a technique to leverage multiple AFS and a TKS as shown in Fig. 1 to implement Anomaly Detection and Correction (ADC) for these jumps.

Jumps in AFS phase or frequency are transferred directly onto the system clock after the loop PLL time constant has elapsed. Over the course of the typical one-day interval between ground updates in a GNSS system small frequency steps of 2E-13 as presented in Formichella (2016) can accumulate to a significant range error of 5.2 m at the end of a day. With the inclusion of at least three operational onboard clocks and a Timekeeping System capable of making rapid and accurate phase measurements of all available clocks (as shown in Fig. 1), it is possible to compare them against each other and positively detect and attribute anomalous behavior in both phase and frequency.

The algorithm is designed to integrate ADC into the FSW ahead of the ensembling, acting on the phase measurements from the PPMs for the three AFS. If an anomaly is detected on any clock, the break is estimated and removed from that clock's raw phase measurement stream. The corrected phase measurements are ensembled and the corrected ensemble phase is then used to calculate the phase error input for the PLL.

Decision logic is used to determine whether any AFS is exhibiting an anomaly and to categorize it as a phase or frequency break. Anomaly detection is performed by comparing the change in each of the AFS phase differences over a range of averaging times. Detection is relative to a statistical noise threshold that becomes more sensitive with increasing averaging time. Since AFS, in general, demonstrate white frequency noise over a wide range of averaging times the Allan Deviation of the frequency difference between two clocks and the corresponding anomaly detection threshold averages down as $\sqrt{\tau}$.

The anomaly detection threshold as a function of averaging time is derived from the Allan Deviation of the frequency difference measurement ($\sigma_y(\tau)$) and adjusted based on a desired rate of false detections per detection event (P_{false}). Since the frequency difference shows white noise, the probability of noise producing a false detection for a given check against a threshold is given as

$$P_{false} = 1 - erf\left(\frac{thresh(\tau)}{\sigma_y(\tau)\sqrt{2}}\right) \tag{2}$$

A family of potential thresholds for different probabilities of false detections are shown in Fig. 4 as a function of the averaging time. Depending on mission needs, the number of averaging times that are checked, and the rate at which a detection check is performed, the required false detection probability per check can be established. Anomaly correction accuracy is given directly by the Allan Deviation of the frequency difference measurement (sigma12 in Fig. 4) and improves for longer averaging times.

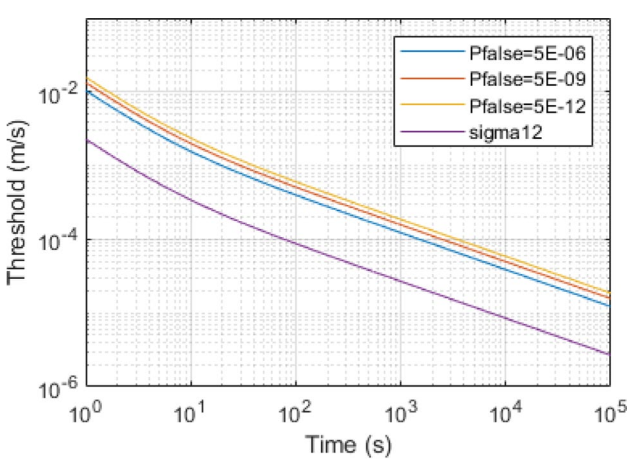

Fig. 4 Detection threshold versus averaging time

Advantages of multiple clocks

The simple weighted average ensembling described in the previous section provides significant advantages when applied to GNSS. The first and most obvious advantage is the enhanced stability of an ensemble. For an ensemble of independent clocks weighted as given in (1), the Allan Deviation is given by

$$\sigma_{en}(\tau) = \sqrt{\left(w_1\sigma_1(\tau)\right)^2 + \left(w_2\sigma_2(\tau)\right)^2 + \left(w_2\sigma_2(\tau)\right)^2} \tag{3}$$

which can be generalized to

$$\sigma_{en}(\tau) = \sigma_{AFS}(\tau)/\sqrt{n} \tag{4}$$

for the special condition where there are n equally weighted clocks, each with identical stability as a function of averaging time. Thus, the stability improves by $\sqrt{n}$ for averaging times τ over which the constituent AFS have identical stabilities. As discussed previously, this grants a 40% improvement in stability with three AFS and a corresponding 40% reduction in the clock contribution to URE.

Another way in which ensembling provides significant improvements for GNSS is when one of the onboard clocks experiences a phase or frequency anomaly. Anomalous phase or frequency excursions that are not the result of stochastic noise processes, and thus gain the full advantage of the weighted average performed by the ensemble. Thus, for an ensemble of n identical AFS in which one of them experiences a frequency anomaly of magnitude Δf,

$$\Delta f_{en} = \Delta f/n \tag{5}$$

such that the impact to the ensemble is reduced by a factor of n.

It has been shown in Oleksak (2021) that a common cause of apparent random walk in GNSS RAFS is infrequent distinct frequency jumps caused by lamp light shifts. Since this noise process is the result of frequency jumps, there will be a factor of n stability improvement for an ensemble over averaging times dominated by this noise process. For GNSS RAFS this is generally over averaging times greater than one day. Finally, the benefits of anomaly correction are self-evident and have been previously presented in Weiss (2010). The correction accuracy for anomaly detection is on the order of the stability of the AFS over the averaging time used to detect the anomaly. This can reduce the range error at the end of the day from meters to centimeters. A summary of the user range error advantages for both ensembling and anomaly correction are summarized in Table 2, where Tdev is the time deviation plus any residual deterministic offset from the anomaly.

While ensembling and anomaly detection and correction cannot fully recover the performance prior to an anomaly,

Table 2 URE Advantages of Ensembling

Parameter	Tdev @ 1 Day (s)	URE @ 1 Day (m)
Single AFS	2.3E-10	0.069
3 AFS Ensemble	1.3E-10	0.040
2E-13 Frequency Anomaly		
Single AFS	2.6E-8	5.2
3 AFS Ensemble	8.6E-9	1.7
Anomaly Correction	2.2E-9	0.67

together they do offer an order of magnitude improvement in URE even for extremely small frequency jumps. For larger, truly anomalous events the performance of the anomaly correction will be similar, while the impact of the anomaly could be an order of magnitude worse.

We have shown that the utilization of multiple AFS on GNSS space vehicles leads to significant performance improvements in clock stability, URE, and anomaly resiliency. Given the existing common practice of equipping GNSS with multiple AFS these advantages can be realized with little additional investment.

References

Camparo J., Formichella V., Sesia I., Signorile G., Galleani L., Tavella P. (2016) Characterization and Consequences of GPS Rb Clock Lamp. Proc. 2016 International Frequency Control Symposium. IEEE Press: Piscataway, NJ

Formichella V (2016) The Allan variance in the presence of a compound Poisson process modelling clock frequency jumps. Metrologia 53:1346

Janis J., Jones M. (2017) Algorithm for On-Orbit GPS III Clock Correction. Proceedings of the 48th Annual Precise Time and Time Interval Systems and Applications Meeting, January 2017, p177–182

Oleksak K., Camparo J. (2021) Clock Ensembling to Mitigate GNSS Atomic Clock Frequency Jumps, Proceedings of the 51st Annual Precise Time and Time Interval Systems and Applications Meeting, January 2021

Taylor C. (2017) Excelitas Technologies Celebrates 20 Years of Continuous, Failure-Free Operations of its Rubidium Atomic Frequency Standard aboard on-orbit GPS IIR Satellite. Press Release, Excelitas Technologies, Waltham, MA, August 11

Weiss M., Shome P., Beard R. (2010) Onboard GPS Clock Monitoring for Signal Integrity, Proceedings of the 42nd Annual Precise Time and Time Interval Systems and Applications Meeting, January 2010, p.465–479

John P. Janis is a Systems Engineering Group Lead at L3Harris Technologies in the Space and Airborne Systems business segment. Dr. Janis is the lead systems engineer for the PNT Division timing systems. Dr. Janis received his Ph.D. in physics from the University of Rochester in 2004 in Atomic Physics and has led PNT timing systems since 2014.

Michael R. Jones is a Lead Systems Engineer at L3Harris Technologies in the Space and Airborne Systems PNT Division. Dr. Jones specializes in all aspects of engineering for the GPS Timekeeping System and Atomic Frequency Standards. Dr. Jones received his Ph.D. in astronomy from Rice University in 1993.

Nicholas F. Quackenbush is a Systems Engineer L3Harris Technologies in the Space and Airborne Systems PNT Division. Dr. Quackenbush specializes in the Timekeeping System for GPS-related systems. Dr. Quackenbush received his Ph.D. in physics from Binghamton University in 2016.