ORIGINAL ARTICLE

Design of a high-stability heterogeneous clock system for small satellites in LEO

Damon Van Buren[1] · Penina Axelrad[1] · Scott Palo[1]

Abstract

We describe our investigation into the performance of low-power heterogeneous timing systems for small satellites, using real GPS observables from the GRACE Follow-On mission. Small satellites have become capable platforms for a wide range of commercial, scientific and defense missions, but they are still unable to meet the needs of missions that require precise timing, on the order of a few nanoseconds. Improved low-power onboard clocks would make small satellites a viable option for even more missions, enabling radio aperture interferometry, improved radio occultation measurements, high altitude GPS navigation, and GPS augmentation missions, among others. One approach for providing improved small satellite timekeeping is to combine a heterogeneous group of oscillators, each of which provides the best stability over a different time frame. A hardware architecture that uses a single-crystal oscillator, one or more Chip Scale Atomic Clocks (CSACs) and the reference time from a GPS receiver is presented. The clocks each contribute stability over a subset of timeframes, resulting in excellent overall system stability for timeframes ranging from less than a second to several days. A Kalman filter is used to estimate the long-term errors of the CSACs based on the CSAC-GPS time difference, and the improved CSAC time is used to discipline the crystal oscillator, which provides the high-stability reference clock for the small satellite. Simulations using GRACE-FO observations show time error standard deviations for the system range from 2.3 ns down to 1.3 ns for the clock system, depending on how many CSACs are used. The results provide insight into the timing performance which could be achieved on small LEO spacecraft by a low power timing system.

Keywords GPS satellite clocks · Allan variance · Kalman filter · Phase locked loop · Small satellite · High-stability clock

Introduction

With low size, weight and power (SWAP) and low launch cost, small satellites are appealing for a variety of earth science and defense missions. Advances in sensors and communication hardware have enabled small satellites to produce imagery and other data which rival much larger spacecraft (Norton et al. 2016; Alminde et al. 2009; Weston et al. 2020; Shah and Arshad 2013). Cubesats have also been sent on deep space missions for the first time with the launch of the MarCo spacecraft to Mars in May, 2018 (Tkacenko et al. 2019).

Accurate position and time estimates are important for many missions, for example, deep space missions using one-way ranging (Imken et al. 2017; Segret et al. 2018), constellations functioning as sensor arrays for astronomy (Rajan et al. 2016), LEO constellations for GPS augmentation and improved positioning (Reid et al. 2018; Ge et al. 2018), and earth observation missions which require accurate position and time information for imagery and science data.

Constellations have been proposed for use as distributed arrays for ultra-low frequency radio astronomy (Rajan et al. 2013). For these missions, clock stabilities are critical to achieving performance requirements by aligning the timing of the constituent satellites. Allan deviations better than 10E−8 are needed for one second integration times, and deviations below 10E−12 are necessary to achieve integration times up to 1E4 s for RF frequencies up to 50 MHz.

The use of broadband commercial LEO constellations to generate GNSS signals has been analyzed (Reid et al. 2018) to identify the requirements on the LEO timing and position errors. Because the commercial constellations have higher numbers of spacecraft, and because LEO offers better

✉ Damon Van Buren
damon@spacesigsci.com

1 Ann and H.J. Smead Aerospace Engineering Sciences, University of Colorado, Boulder, CO 80309, USA

geometry for navigation, time and position error requirements are relaxed by a factor four compared to GPS MEO satellites (3.0 m for LEO vs. 0.8 m for GPS). To achieve timing accuracy comparable to GPS, these studies recommend using Chip Scale Atomic Clocks (CSACs) with time error updates once per orbit.

For small satellites to serve in these additional important roles, they must meet more challenging system requirements. However, small satellites present a challenging environment for improvement for several reasons. They have much lower volume and power than larger satellites, so any system improvements must occur within the limited SWAP envelope. Small satellites also do not have the ability to provide temperature control for sensitive components such as clocks. And because cost is a big part of why small satellites are appealing, improved systems must still fit within the reduced budget.

Heterogeneous timing system

Heterogeneous clock systems offer a solution that meets the need of small satellites for low SWAP and cost while still providing accuracy on the order of a few ns. By efficiently combining low power clocks having complimentary stability over various timeframes, excellent overall stability may be achieved with lower power and less expense. The timing system outputs a high-stability reference clock, which is used by other payload and communication systems on the spacecraft. This research investigates the integration of several different clock types to provide high performance timekeeping onboard LEO spacecraft.

The key elements of the system are low-power clocks or timing sources with complementary stability, algorithms to optimally combine the clocks, and an efficient, low power hardware architecture, as outlined in the following sections.

System clocks

Since the goal is to provide accurate timing for small satellites while using only a few Watts, constituent clocks are selected based on two criteria: (1) stability over a desired subset of averaging times, and (2) low power utilization. The systems in this study use a single-crystal oscillator (XO), one or more Chip Scale Atomic Clocks (CSACs), and the reference time from a GPS receiver. These clocks each contribute stability over a subset of timeframes, resulting in excellent system stability for timeframes ranging from less than a second to several days.

CSACs were introduced in 2011 (Lutwak 2011; Micro-Semi 2019). These devices provide very good stability for timeframes of a few hours (Cash et al. 2018; Fernández et al. 2017), with very low power (1/8 W) and a small form-factor,

and they are relatively inexpensive (< \$10 k). However, CSACs alone do not provide sufficient long-term stability without external synchronization (Weaver et al. 2012; Bagala et al. 2016a, 2016b). The original SA.45 MicroSemi CSACs also had relatively poor stability for timeframes less than 100 s, and have much higher phase noise than less expensive crystal oscillators. The phase noise issue has since been addressed by the introduction of a low-noise version of the CSAC, but it requires higher power, 295 mW, and is slightly larger (MicroChip 2019).

Because CSACs require so little power, more than one may be included in the heterogeneous system to provide improved overall stability. The stability provided by multiple combined CSACs improves by the square root of the number included ($\sqrt{N}$), so four CSACs are twice as stable as a single CSAC, and sixteen CSACs offer a $4\times$ improvement in stability while requiring only 2 W of total power. Note that these stability improvements do not apply to shared environmental impairments such as thermal variation (Rybak et al. 2018).

GPS is unsurpassed for long-term time stability in non-laboratory environments due to the regular updates to the estimated GPS satellite ephemerides and time errors. GPS has become the de facto standard for orbit estimation for LEO spacecraft, and GPS receivers are often used to discipline other oscillators to improve long-term timing stability in terrestrial applications. GPS offers a straightforward solution to the problem of estimating long-term system time errors and is continuously available for spacecraft in LEO orbits. It is also low SWAP; dual-frequency commercial GPS receivers require approximately 2 W, and are small and lightweight (NovAtel OEM719 Multi-Frequency GNSS Datasheet). Not all GPS receivers support operation in LEO, however.

Correctly modeling GPS receiver timing errors on LEO spacecraft can be nontrivial, however. The dynamic nature of LEO orbits results in rapid changes in signal strength and environmental factors such as the ionosphere. Signal in Space Ranging Errors (SISRE), caused by errors in the predicted GPS orbits and clocks, result in jumps in time and position estimates onboard LEO spacecraft as the visibility of the GPS constellation changes. Prior work investigating GPS errors on LEO spacecraft in detail found that SISRE errors are dominant for dual-frequency receivers (Van Buren et al. 2019). To capture realistic on-orbit errors, we use 10 days of observations from the GRACE Follow-On (FO) mission to calculate GPS receiver time errors, which are representative of what will be seen by other LEO spacecraft.

Two MicroSemi crystal oscillators were chosen to provide short-term stability to the clock system, an Evacuated Miniature Crystal Oscillator (EMXO) and an Oven Controlled Crystal Oscillator (OCXO). The MicroSemi EX-421 is an EMXO that provides excellent stability for one second averaging times and requires only 0.25 W. The OX-174 is

an OCXO from MicroSemi which offers exceptional stability for timeframes from 1 to 1000 s, but requires 1.8 W of continuous power due to the oven.

Figure 1 shows the Allan deviations for the clocks used in the study. The EMXO crystal oscillator is most stable over 1 s intervals, with a minimum ADEV of 2E−11. The OCXO provides excellent stability for timeframes from 0.1 s to 1000 s with ADEVs below 5E−12. Typical ADEV curves for one, four, and sixteen combined CSACs are also shown, with minimum ADEVs occurring at 10,000 s (Lutwak 2011). The GPS clock error was derived from point solution calculations with the GRACE-FO P12 ionospheric-free pseudoranges. The GPS ADEV is worse than the CSACs and the XOs for short timeframes but continues to improve for longer timeframes as expected.

Clock models

Simulated errors for the clocks were generated using state-space models, aside from the GPS timing error, which was calculated from the GRACE-FO observations. The power law clock model includes five elements: Random Walk FM (RWFM), Flicker Frequency Modulation (FFM), White Frequency Modulation (WFM), Flicker Phase Modulation (FPM), and White Phase Modulation (Ramian 2015; Chang et al. 2004). State space clock models are created using a few of the power law components, and good results can often be obtained by using only the WFM and RWFM elements of the power law model (Galleani 2008). For a clock with WFM and RWFM, the clock state propagates according to the following equations, with τ representing the sample period:

Fig. 1 Member clock Allan deviations for N CSACs, EMXO and OCXO, and GPS user clock estimates

$$x_{k+1} = x_k + \tau y_k + \varepsilon_k \tag{1}$$

$$y_{k+1} = y_k + \eta_k \tag{2}$$

$$\varepsilon_k \sim N\left(0, \sigma_{\mathrm{WFM}}\right) \tag{3}$$

$$\eta_k \sim N\left(0, \sigma_{\mathrm{RWFM}}\right) \tag{4}$$

where x represents the clock phase and y represents the clock frequency. The values ε and η represent the process noise in the clock states. The corresponding state propagation matrix is:

$$\Phi = \begin{pmatrix} 1 & \tau \\ 0 & 1 \end{pmatrix} \tag{5}$$

The process noise covariance matrix Q is given by:

$$Q = \begin{pmatrix} \left[\sigma_{\mathrm{WFM}}^2\tau + \frac{\sigma_{\mathrm{RWFM}}^2\tau^3}{3}\right] & \left[\frac{\sigma_{\mathrm{RWFM}}^2\tau^2}{2}\right] \\ \left[\frac{\sigma_{\mathrm{RWFM}}^2\tau^2}{2}\right] & \left[\sigma_{\mathrm{RWFM}}^2\tau\right] \end{pmatrix} \tag{6}$$

Because the OCXO has a large region of relatively flat stability between 0.1 s and 1000 s, three Markov processes were added to the state propagation matrix and the Q matrices to model the FFM errors, as outlined by Davis et al. (2005). The Markov process is defined by the following differential equations, where $1/R$ is the time constant, and $n(t)$ is a continuous white noise process with deviation M:

$$\frac{\mathrm{d}x}{\mathrm{d}t} = y(t) \tag{7}$$

$$\frac{\mathrm{d}y}{\mathrm{d}t} = -Ry(t) + n(t) \tag{8}$$

The state propagation matrix and process covariance matrices for y are then:

$$\Phi = \begin{pmatrix} 1 & \frac{1-\exp(-R\tau)}{R} \\ 0 & \exp(-R\tau) \end{pmatrix} \tag{9}$$

$$Q = \sigma_M^2 \begin{pmatrix} \tau^3 a_{11}(R\tau) & \tau^2 a_{12}(R\tau) \\ \tau^2 a_{12}(R\tau) & \tau a_{22}(R\tau) \end{pmatrix} \tag{10}$$

where

$$a_{11}(x) = \frac{-3/2 + x + 2\exp(-x) - \exp(-2x)/2}{x^3} \tag{11}$$

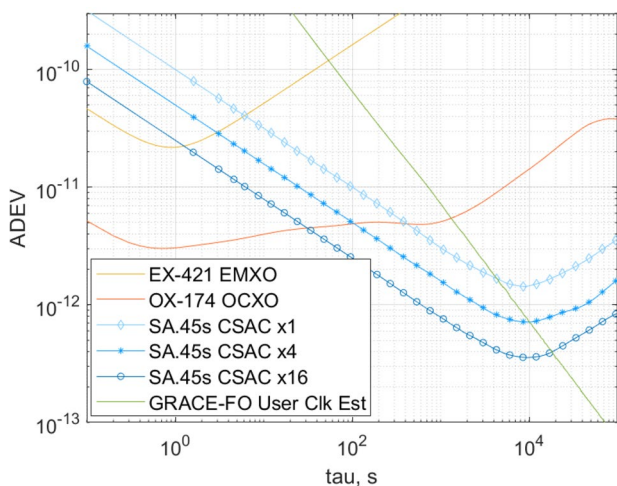$$a_{12}(x) = \frac{1/2 - \exp(-x) + \exp(-2x)/2}{x^2} \tag{12}$$

$$a_{22}(x) = \frac{1 - \exp(-2x)}{2x} \tag{13}$$

The combined clock model which includes WFM, RWRM and three Markov processes is then:

$$\Phi_{ii} = \begin{pmatrix} 1 & \tau_0 & \frac{1-\exp(-R_1\tau_0)}{R_1} & \frac{1-\exp(-R_2\tau_0)}{R_2} & \frac{1-\exp(-R_3\tau_0)}{R_3} \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \exp(-R_1\tau_0) & 0 & 0 \\ 0 & 0 & 0 & \exp(-R_2\tau_0) & 0 \\ 0 & 0 & 0 & 0 & \exp(-R_3\tau_0) \end{pmatrix} \tag{14}$$

where Q_{ii} is defined as:

Fig. 2 ADEV for a notional 3-clock ensemble

$$Q_{ii} = \begin{pmatrix} Q_{ii,11} & \frac{\sigma_{RWFM}^2 \tau_0^2}{2} & \tau_0^2 a_{12}(R_1\tau_0)\sigma_{M1}^2 & \cdots & \tau_0^2 a_{12}(R_3\tau_0)\sigma_{M3}^2 \\ \frac{\sigma_{RWFM}^2 \tau_0^2}{2} & \sigma_{RWFM}^2 \tau_0 & 0 & 0 & 0 \\ \tau_0^2 a_{12}(R_1\tau_0)\sigma_{M1}^2 & 0 & \tau_0^2 a_{22}(R_1\tau_0)\sigma_{M1}^2 & 0 & 0 \\ \vdots & 0 & 0 & \ddots & 0 \\ \tau_0^2 a_{12}(R_3\tau_0)\sigma_{M3}^2 & 0 & 0 & 0 & \tau_0^2 a_{22}(R_3\tau_0)\sigma_{M3}^2 \end{pmatrix} \tag{15}$$

and $Q_{ii,11}$ is defined as:

$$Q_{ii,11} = \sigma_{WFM}^2 \tau_0 + \frac{\sigma_{RWFM}^2 \tau_0^3}{3} + \tau_0^3 a_{11}(R_1\tau_0)\sigma_{M1}^2 + \tau_0^3 a_{11}(R_2\tau_0)\sigma_{M2}^2 + \tau_0^3 a_{11}(R_3\tau_0)\sigma_{M3}^2 \tag{16}$$

Note that the OCXO model, which included FFM processes, was used only to generate simulated OCXO clock errors.

Methods for combining clocks

The two most common methods for combining clocks are disciplining using Phase Locked Loops (PLLs), and timescale generation using Kalman filters. The combined clock may be physical, as is the case when a GPS receiver disciplines a crystal oscillator to minimize long-term errors (Lombardi 2008), or virtual, as is the case when generating a timescale from a clock ensemble in a lab setting (Greenhall 2012; Senior and Coleman 2014; Stein 2003). Figure 2 illustrates the ADEVs for a conceptual 3-clock ensemble. The combined clock reflects the performance of the best constituent clock at each timeframe.

PLLs have been used for several decades to discipline crystals with good short-term stability and phase noise with oscillators having better long-term stability and frequency accuracy (Best 2007). Kalman filters have also been used for
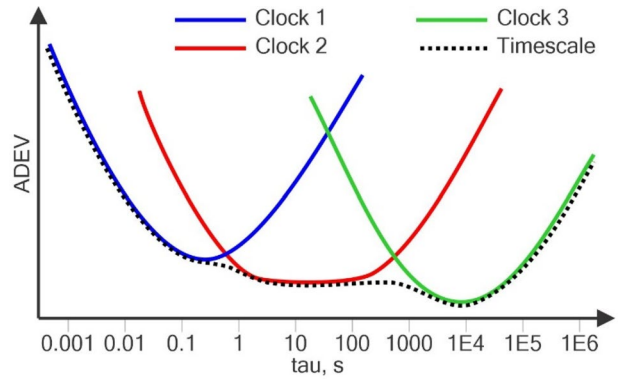
decades to optimally combine atomic clocks in a lab setting (Senior and Coleman 2014).

Kalman filter for combining GPS and CSACs

The heterogeneous system first combines one or more CSACs with a GPS reference time, using a Kalman filter. For systems with multiple CSACs, the CSAC time offsets can be combined optimally by averaging since they share the same power law characteristics. The errors of the GPS reference time have characteristics similar to white phase noise, with a − 1 slope on Allan deviation plots. For this reason, we elect to treat the GPS timing errors as measurement noise, and use a Kalman filter to estimate the states of the combined CSACs.

A Kalman filter is ideal for estimating the long-term CSAC errors from the CSAC-GPS combination (Galleani and Tavella 2003). Since this estimate will be subtracted from the CSAC-EMXO difference, it exists only as a numerical signal value inside the processing hardware. The Kalman filter also lends itself to more complex models of the CSAC errors, such as a sinusoidal temperature induced error (Rybak et al. 2018).

The success of this approach partly hinges on how the Kalman filter will react to the GPS time errors. Kalman filters do an excellent job of estimating system states with white measurement noise, so we would expect the overall clock system performance to reflect this. However, although the ADEVs of the GPS time errors appear to have a -1 slope similar to white phase noise, the GPS time errors are not white phase noise. Instead, they consist of sharp jumps caused by SISRE and GPS visibility changes, connected by slowly varying errors that occur while the GPS visibility remains constant.

The Kalman filter algorithm operates recursively as new observations arrive (Brown and Hwang 2012). First, a new a priori estimate of the state is generated from the a posteriori estimate of the previous state using the state transition matrix:

$$\hat{x}_n^- = \Phi \hat{x}_{n-1}^+ \tag{17}$$

The a priori estimate of the error covariance matrix is generated, also using the state transition matrix and the a posteriori error covariance matrix:

$$P_{Rn}^- = \Phi P_{R(n-1)}^+ \Phi^T + Q \tag{18}$$

The Kalman gain is computed using the a priori error covariance matrix and the observation matrix:

$$K_{Rn} = P_{Rn}^- H^T \left(H P_{Rn}^- H^T + R_{GPS} \right)^{-1} \tag{19}$$

Although we model GPS clock bias errors in the filter assuming white noise with variance R, actual errors seen in real data are more complex (Van Buren et al. 2019). For this reason, the measurement noise parameter R is hand optimized to minimize the time error standard deviation.

Finally, a new a posteriori state estimate and error covariance matrix are generated by optimally combining the a priori estimate with the observation error weighted by the Kalman gain:

$$\hat{x}_n^+ = \hat{x}_n^- + K_{Rn} \left(y_n - H \hat{x}_n^- \right) \tag{20}$$

$$P_{Rn}^+ = \left(I - K_{Rn} H \right) P_{Rn}^- \tag{21}$$

Using this approach, the algorithm updates the state estimate each time a new observation is received. The time between observations need not be constant, provided the state transition matrix and the covariance matrix for the forcing function are computed for the current sample interval.

Phase locked loops

Phase locked loops are commonly used to discipline an oscillator with good short-term stability such as a crystal
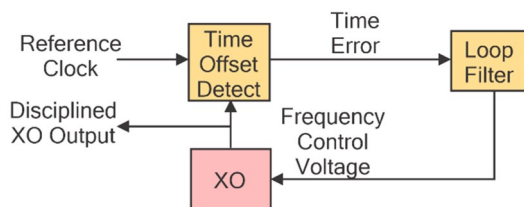

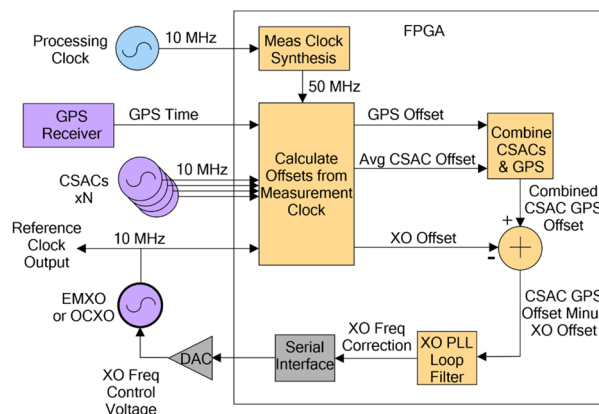**Fig. 3** PLL block diagram

Fig. 4 Generic hardware block diagram showing member clocks and FPGA processing functions

oscillator, using a timing source with better long-term stability, such as GPS. The measured time offset between the crystal and GPS is passed to a loop filter, which generates a frequency offset, which is passed to the frequency control input of the crystal to complete the loop, as shown in Fig. 3. The bandwidth of the loop filter is selected to provide the best combined response, so only the long-term differences between the XO and GPS affect the frequency control voltage (Best 2007). For this research, the difference between the crystal oscillator and the combined timing of the CSACs and GPS is used as the time error.

Small satellite clock system architecture

The processing hardware for the heterogeneous system must provide the following functions:

- Measure the time offsets of the member clocks
- Implement algorithms to combine clocks
- Generate a frequency control voltage to steer a reference clock output

Ideally, the processing hardware should be low power and low volume while still providing enough processing capability to implement the necessary clock combining algorithms.

Figure 4 shows the hardware architecture of the notional system. A Field Programmable Gate Array (FPGA) was selected for implementing the algorithms. FPGAs are more power efficient than general purpose processors, are reprogrammable, and have flexible IOs which may be configured to interface with a variety of digital signal types.

N CSAC clocks are connected to the FPGA, along with a single-crystal clock and a GPS time reference. To keep the system design as simple and low-power as possible, square wave CMOS clock signals are connected directly to FPGA input pins, rather than sampling the clocks at a high rate with an Analog to Digital Converter (ADC). The FPGA receives 10 MHz CMOS signals from the clocks and the time reference from the GPS receiver on several of its input pins, as well as a 10 MHz CMOS processing clock. It outputs a serial data stream to a high-precision, low rate DAC, which generates the frequency control voltage for a crystal oscillator.

The FPGA's internal clock management circuitry is used to synthesize a measurement clock at a frequency near 50 MHz. The clock management circuitry can synthesize clocks with an K/L ratio to the reference clock, where K and L are integers and are reprogrammable, enabling us to fine-tune the measurement clock frequency. The offset of each member clock from the measurement clock is measured and passed to the clock combining algorithms.

Since the N CSACs have the same stability characteristics, their offsets are averaged to provide a $\sqrt{N}$ improvement in stability compared to a single CSAC. The CSAC average is then combined with the GPS receiver time using a Kalman filter to produce a timing estimate with better long-term stability.

The difference between the combined N-CSAC + GPS clock and the crystal oscillator is then calculated and passed to a loop filter. The output of the loop filter is the long-term frequency error of the XO relative to the N-CSAC + GPS clock. The frequency correction from the loop filter is sent out of the FPGA over a serial interface to the high precision DAC, which generates an analog frequency control voltage for the XO to discipline it to the N-CSAC + GPS combination.

Fine clock offset measurements

A variety of methods may be used to detect the time differences between member clocks, depending on the clock signal type. For example, sinusoidal member clocks may be sampled with Analog to Digital Convertors (ADCs) and signal processing used to estimate the phase offset between them.

For this work, we use a common high-frequency measurement clock to observe the member clock offsets by counting the number of measurement clock cycles prior to the rising edges of each member clock. As the member clocks vary over time due to clock errors, the number of measurement clocks prior to the rising edges of the member clocks changes, indicating how much the member clock times have moved relative to the common measurement clock.

This approach has the benefit of being extremely simple to implement. Member clocks with CMOS square wave outputs are used, so each member clock output may be connected directly to a single FPGA input pin. The clock signals are double registered in the programmable logic of the FPGA using the measurement clock, and a rising edge detector is used to register the measurement clock count at each rising edge of the member clock. The difference between the expected measurement clock count for a given number of member clock cycles and the actual count is the observed member clock offset.

Because the algorithms for combining the clocks rely on the differences in the member clock offsets, the precision of the clock offset measurements has a significant impact on the system performance. Most FPGAs, even those with limited processing capability, are able to synthesize higher frequency clocks from a lower frequency input clock, and we can use this feature to synthesize the high-frequency measurement clock using the FPGA clock management resources. However, if the measurement clock frequency is an integer multiple of the 10 MHz member clock frequency, the precision of the offset measurement approach is limited to the period of the measurement clock. For low-performance FPGAs, the maximum synthesizable frequency is several hundred MHz, limiting the time offset measurement precision to several ns.

We can improve the time offset measurement resolution by making the measurement clock frequency slightly higher or lower than an integer multiple of the member clock frequency, so the measurement clock "walks" relative to the member clocks. For example, suppose the measurement clock frequency is chosen so that it is slightly higher than an integer multiple of the member clock. In that case, the measurement clock cycles will slowly advance relative to the member clock cycles over time. This will also cause the number of measurement clock cycles between rising edges of the member clock to vary by one from a nominal value occasionally. For example, when using a measurement clock with a frequency a little above 20 MHz, the number of measurement cycles per 10 MHz member clock will vary by one, producing a series of two-cycle counts broken up by an occasional three-cycle count.

Changes in the relative time offset of the 10 MHz member clocks will cause advances or delays of the three-cycle count. For example, if the 10 MHz CSAC clock is slightly delayed (slow), the three-cycle measurement clock count will occur earlier. By keeping track of the number of member clock cycles between the three-cycle counts of the measurement clock, we can detect relative changes in the member
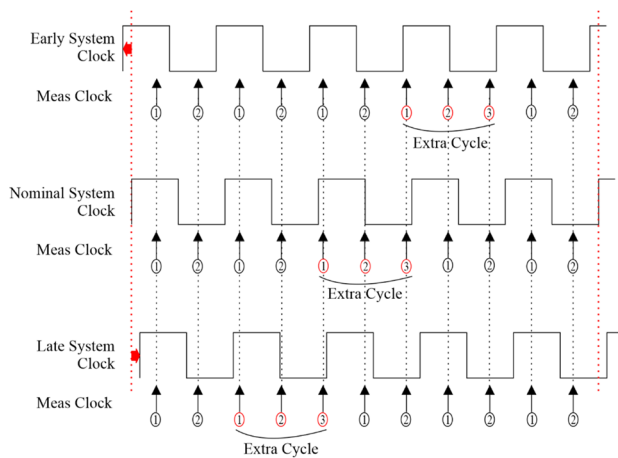
Fig. 5 Walking measurement clock versus member clock

Fig. 6 Fine offset measurement simulation for 78 ps resolution, equivalent to 12.8 GHz measurement clock frequency

clocks with much higher precision than the measurement clock period. These changes are accumulated to track the total offset of each of the member clocks.

Figure 5 shows the measurement clock counts per member clock cycle for two early, nominal, and late member clocks. In this case, there are nine measurement clocks for every four cycles of the 10 MHz member clock, so the measurement clock frequency is 22.5 MHz. It is apparent that although the early and late member clocks are only shifted by a fraction of the 22.5 MHz measurement clock period, the shifts change the timing of the occurrence of the three-cycle measurement clock. In the case when the member clock is slightly early relative to the measurement clock, the three-cycle measurement occurs one member clock cycle later, or five-member clock cycles after the previous three-cycle instance instead of the usual four. For the case when the member clock is delayed slightly, the three-cycle measurement clock count occurs one member clock cycle earlier, or three-member clock cycles after the previous instance instead of four.

This demonstrates that the approach is providing time offset resolution, which is smaller than the period of the measurement clock. The time resolution is the reciprocal of the product of the measurement clock frequency and the number of member clock cycles required for the measurement clock to add one extra cycle. In this case, the resolution is 1/(22.5E6*4), or 11.1 ns, which is equivalent to using a 90 MHz measurement clock. The member clock time offset estimate is updated by adding or subtracting 11.1 ns when the interval between the three-cycle measurement clock counts is three- or five-member clock cycles. If the interval is four-member clock cycles, the time offset estimate remains the same.

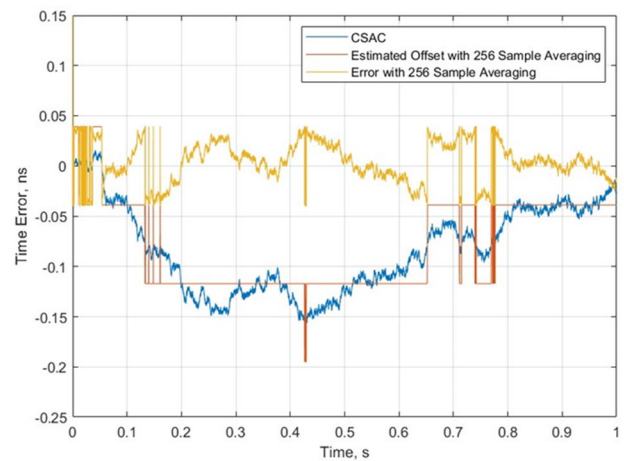We can improve the resolution of the time offset measurements by either increasing the measurement clock frequency

or moving the measurement clock frequency closer to an integer multiple of the member clock frequency. The latter case is equivalent to taking more member clock cycles to add a single measurement clock cycle. For this research, we have chosen a measurement clock of 50.039 MHz, so there are 5.0039 measurement clock cycles per member clock cycle. The measurement clock completes one extra cycle for every 256 cycles of a 10 MHz member clock, or 1281 cycles instead of 1280. For each cycle that the interval between the six cycle occurrences differs from the nominal 256 member clock cycles, the estimated offset of the member clock is changed by 1/(50.039E6*256), or 0.078 ns. This is equivalent to a measurement clock frequency of 12.82 GHz—a frequency well beyond the capability of even high-performance FPGAs.

Because updates in the time offset measurements only occur every 256 cycles of the member clock, it is important to consider the impact of this added latency on the clock combination algorithms. The clock combination algorithms operate by estimating and removing the long-term variation of the member clocks. For the CSAC-GPS combination, the CSAC errors for averaging times longer than 1000 s are estimated. For the EMXO and OCXO disciplining via PLL, errors for averaging times longer than 10 s are tracked. Since the added latency due to waiting 256 10 MHz clock cycles is 256/10E6, or 25.6 us, we can safely assume that the additional delay will have no meaningful impact on the clock combination algorithms, while providing a dramatic improvement in the time offset resolution.

It is also important to note that the quality of the measurement clock has no significant impact on the performance of the offset measurements. Because the measurements are nearly simultaneous, the measurement clock error is common to all member clock offset observations. The Kalman

filter and the PLL operate on time differences between the member clocks, so the measurement clock variations are removed.

Figure 6 shows the simulated results of using this technique. The offset of a single CSAC signal is quantized using a 50.04 MHz measurement clock, and the offsets are updated every 256 cycles of the CSAC. The quantization step size is 1/(50.04E6 * 256), or 0.08 ns, limiting the measurement error to ± 0.04 ns.

Heterogeneous clock system with GRACE-FO observations

For this case study we used real Gravity Recovery and Climate Experiment Follow-On (GRACE-FO) mission (Kornfeld et al. 2019) observations to simulate the GPS time errors in the heterogeneous system. GRACE-FO observations include receiver noise with varying SNR due to elevation changes, delays in GPS signal acquisition, and carrier phase jumps, providing realistic GPS time errors for our simulations, and the GRACE-FO orbit has been estimated with an accuracy of several cm.

The GRACE-FO mission continues the highly successful mission of the original GRACE gravity mapping mission, and is a joint venture between NASA and the German Research Centre for Geosciences. The two GRACE-FO spacecraft are separated by an average of 220 km, and are placed in a near-circular polar orbit with an altitude of 490 km. Changes in the distance between the satellites are measured precisely using a microwave link with Dual One-Way Ranging (DOWR) which is accurate to a few microns and used to estimate earth's gravity.

NASA operates the Physical Oceanography Distributed Active Archive Center (PODAAC) website (NASA 2021), which hosts the GRACE-FO GPS observation data and post-processed orbit data. The GRACE-FO satellites are equipped with dual-frequency BlackJack GPS receivers (Montenbruck and Kroes 2003), which capture observations of pseudorange and carrier phase on the L1 and L2 frequencies at 10 s intervals. For this research, the pseudorange observables P1 and P2 were used to form ionosphere free pseudoranges observations P12, and the phase observables L1 and L2 were used to form ionospheric-free phase observations L12.

GRACE-FO GPS antenna position and clock offset compensation

In order to use the GRACE-FO observations in the clock simulations two offsets were accounted for. The first is the offset of the GPS antenna from the spacecraft center of mass. The primary GPS antenna is offset in the zenith and along-track dimensions from the center of mass, while the JPL

post-processed orbit estimates provide the position of the spacecraft center of mass. A radial offset of 0.48 m and an along-track offset of 0.3 m (Wen et al. 2019) were subtracted from the position estimates calculated from the pseudoranges to reflect the GPS antenna position on the spacecraft relative to the spacecraft center of mass.

To realistically simulate the LEO GPS receiver clock errors, the GRACE-FO user clock error was also estimated and removed. Expected pseudoranges were calculated using post-processed estimates of GPS antenna phase center positions and time errors from the National Geospatial Intelligence Agency (NGA), along with post-processed GRACE-FO position estimates from JPL (Wen et al. 2019). This offset was removed from the user clock estimates calculated in the point solutions using GRACE-FO observed pseudoranges to provide realistic LEO GPS receiver time errors for the time system simulations.

Simulations

Simulations of the heterogeneous clock system covering a 10-day period were run to demonstrate performance and the impact of important design choices such as time offset measurement precision and the number of CSACs used. GRACE-FO observations from November 16–25 of 2018 were used to generate ionospheric-free carrier phase and pseudorange values, from which point solutions of position and time were calculated. Simulated CSAC and XO clocks were generated using state-space models described above. All reported ADEVs, standard deviations, and maximum errors cover the same 10-day timeframe.

The GRACE-FO observations also included carrier phase observations L1 and L2, enabling the use of a Hatch filter for carrier smoothing of pseudoranges. The P12 and Hatch filtered GPS time ADEVs are shown in Fig. 7, along with
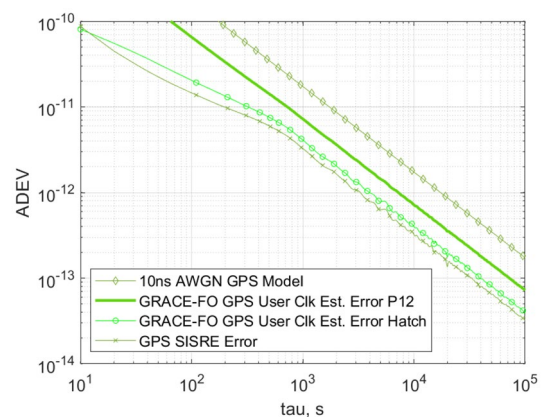

Fig. 7 Allan deviations of GPS time estimate errors using GRACE-FO observations

the GPS time error from just the SISRE for GRACE-FO from the same time period. For reference, the ADEV from the 10 ns AWGN signal which was previously used to model the GPS time error, is also included.

The time error caused by only SISRE is the lowest, with the time error of the point solutions computed with carrier smoothed pseudoranges using the Hatch filter slightly higher. The time error of the point solutions computed with only the P12 ionospheric-free GRACE-FO pseudoranges is higher than the SISRE and carrier smoothed errors, but still lower than the 10 ns AWGN model.

Standard deviations of the various time errors are given in Table 1. Point solutions were calculated using GRACE-FO pseudoranges only (P12), as well as pseudoranges smoothed using a Hatch filter. Simulated pseudorange observations which include only estimated SISRE were also used. The standard deviation of the GRACE-FO time estimate is very close to the simulated SISRE-only time error, slightly higher than 2 ns. The standard deviation of the time error using ionospheric-free pseudoranges is over 4 ns. Although the Hatch filtered pseudoranges provide lower time errors, they do not result in better system performance, as we will see in the remaining system simulations. Note that the estimated GRACE-FO clock offsets were removed from the point solutions using GRACE-FO P12 and Hatch filtered pseudoranges.

CSAC GPS combined clock

Figure 8 shows the ADEVs for CSACs combined with GPS time estimates using the P12 observations, using a Kalman filter. The process noise matrix Q has been computed to accurately model the performance of the number of included CSACs. The measurement covariance R has been adjusted by hand to minimize the standard deviation of the time error for one and sixteen CSACs. The Kalman filter combines the two clock types very effectively, resulting in an improvement in stability for timeframes above 1000 s over the GPS time estimate with P12. The system with sixteen CSACs is more stable over the entire range of timeframes than the single-CSAC system.

It is interesting to note the improvement in the overall stability for timeframes greater than 1000 s, showing dramatic
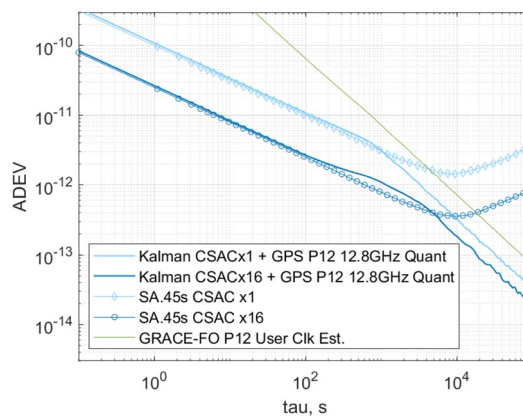
Fig. 8 ADEVs of Kalman filter combination of CSACs + GRACE-FO GPS P12

improvement provided by the Kalman filter over the GPS time errors. Note that the measurement noise parameter R was hand-optimized in each case to minimize the standard deviation of the overall error for each system. Because the long-term errors are dominant, their reduction results in the improvement of the time error standard deviation of the system.

Table 2 shows the standard deviations and maxima of the time errors for November 16–25 of 2018 after combining CSACs with time estimates from the GPS point solutions, which were calculated using only the P12 pseudorange observations. The time error standard deviation for a single CSAC combined with GPS improves the standard deviation from 4.2 to 1.9 ns. The maximum error is greatly reduced, from 36.6 to 10.5 ns. The sixteen CSAC system further reduces the standard deviation and maximum to 1.1 ns and 5.1 ns, respectively.

For timeframes above 1000 s, the combined N-CSAC + GPS time stability shows an improvement of 2–3 times over the GPS time estimates calculated using P12 pseudoranges. Based on the improvement in ADEVs and standard deviations for GPS time estimates calculated using carrier smoothing, we would expect that a system using carrier smoothed pseudoranges would provide even better time stability for averaging times longer than 1000 s. However,

Table 1 Standard deviations and maxima of GPS time errors

GPS model used	Time error std. dev. ns	Max time error, ns
10 ns AWGN model	10.0	51.3
Point Sol. with GRACE-FO P12	4.2	36.6
Point Sol. with GRACE-FO Hatch	2.3	15.1
Point Sol. with SISRE Only	1.9	13.7

Table 2 Time error standard deviations and maxima of CSACs + GRACE-FO GPS P12

System Type	Time error std. dev. ns	Max time error, ns
GRACE-FO P12 GPS	4.2	36.6
CSACx1 GRACE-FO P12 GPS	1.9	10.5
CSACx4 GRACE-FO P12 GPS	1.5	8.7
CSACx16 GRACE-FO P12 GPS	1.1	5.1

the combined N-CSAC + GPS time estimates with carrier smoothed pseudoranges are slightly worse for timeframes above 1000 s. Since pseudoranges smoothed using the Hatch filter do not offer any improvement over the unsmoothed pseudoranges, the remaining heterogeneous clock system simulations all use the GPS time estimates computed using the P12 ionospheric-free pseudoranges.

Effect of clock offset measurement precision

The precision of the time offset measurement has a significant impact on the overall system performance. Figure 9 shows the ADEVs of the quantization errors for different measurement clock frequencies. The GRACE-FO GPS clock estimate using P12 is included for comparison. Clearly, a time offset measurement frequency of 50 MHz would impair the overall system performance, since the ADEV is higher than the GPS P12 time ADEV.

Figure 10 confirms this by showing the combined CSAC-GPS performance for time offset measurement frequencies of 50 MHz and 12.8 GHz. While the system with a single CSAC shows only slightly degraded performance with the 50 MHz measurement clock, the sixteen CSAC system is significantly impaired by the lower measurement precision, especially for timeframes below 1000 s. This result demonstrates that for high performance systems, high precision measurement of the member clock offsets is critical.

This is an important consideration for spacecraft using commercially available GPS chips. Some devices provide the timing output at a fixed one pulse-per-second (PPS) rate, and with relatively large errors (20 ns). This low sample rate and high maximum error would prevent use of the fine offset measurement technique described previously. Some GPS devices also provide the estimated user clock offset via a software interface, which eliminates the need for measuring the GPS timing signal altogether, provided the measurement
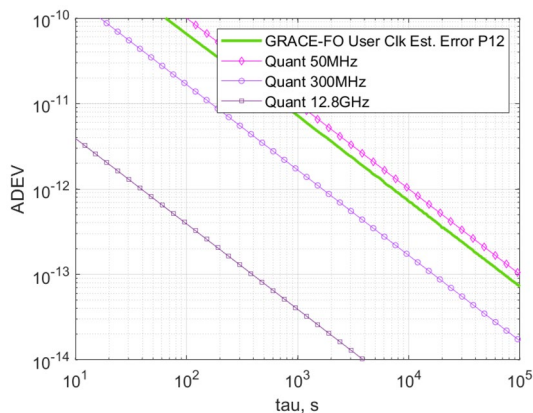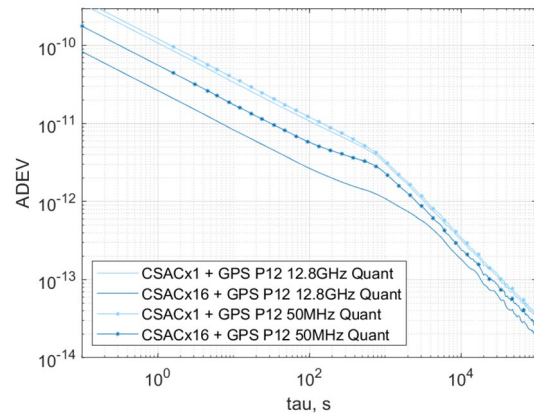

Fig. 10 ADEV comparison of CSACs + GPS with 12.8 GHz time offset measurement versus 50 MHz time offset measurement

clock is used as the reference clock for the GPS device. Although the software interface is slow compared to direct hardware measurements of the GPS time offset, the latency is small compared to the timeframes in which GPS contributes to the combined system clock, which are beyond 1000 s.

Figure 11 shows the combined EMXO-CSAC-GPS results with one and sixteen CSACs. As before, the measurement noise covariance (R) used in the Kalman filter has been hand-optimized to provide the minimum standard deviation of the system time error. The N-CSAC + GPS clocks are differenced with the EMXO, and the estimated EMXO errors are passed through a loop filter, which generates a frequency offset value inside the FPGA. The frequency offsets are passed out of the FPGA to a serial DAC, which generates the analog voltages needed to discipline the EMXO. The PLL bandwidths have been optimized to provide the best combined stability.

Table 3 shows the standard deviations and maxima of the heterogeneous systems including the EMXOs. The EMXO
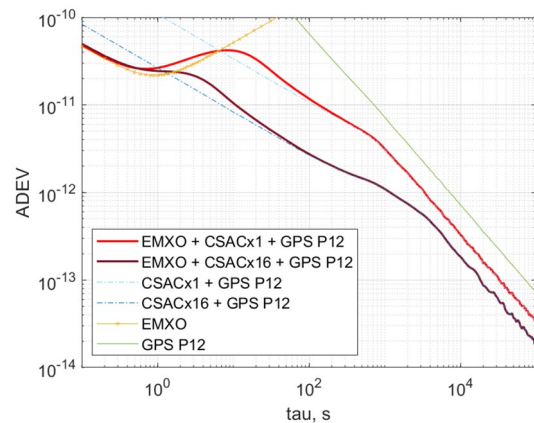

Fig. 9 ADEVs of quantization errors for different time offset measurement resolutions
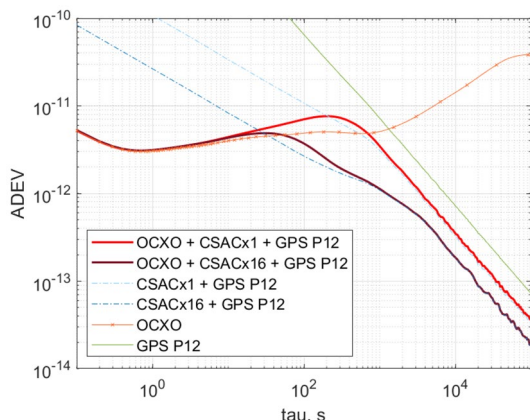

Fig. 11 ADEVs for heterogeneous system with EMXO + CSACs + GPS

Table 3 Time error standard deviations and maxima for EMXO + CSACs + GRACE-FO GPS

System type	Time error std. dev. ns	Max time error, ns
EMXO + CSACx1 + GRACE-FO GPS	1.9	10.5
EMXO + CSACx4 + GRACE-FO GPS	1.5	8.7
EMXO + CSACx16 + GRACE-FO GPS	1.1	5.1

Fig. 12 ADEVs for OCXO + CSACs + GPS

Table 4 Time error standard deviations and maxima for OCXO + CSACs + GRACE-FO GPS

System type	Time error std. dev. ns	Max time error, ns
OCXO + CSACx1 + GRACE-FO GPS	2.1	11.0
OCXO + CSACx4 + GRACE-FO GPS	1.5	8.9
OCXO + CSACx16 + GRACE-FO GPS	1.1	5.1

addition only affects ADEVs below 100 s, but the standard deviations are controlled by the long-term stability, beyond 1000 s. For this reason, the standard deviations and maxima are unchanged vs. the CSAC-GPS systems.

Figure 12 shows the ADEVs for the combined OCXO-CSAC-GPS systems. As with the EMXO, a PLL is used to discipline the OCXO based on the time difference between the OCXO and the CSAC-GPS combination, with the PLL bandwidths optimized to provide the best combined system performance.

Table 4 shows the standard deviations and maxima for the OCXO-CSAC-GPS systems. In this case the PLL disciplining circuit is unable to follow the transition from the CSAC-GPS response to the OCXO response, resulting in slightly worse performance for the system around 300 s, especially for the system with a single CSAC. The PLL bandwidth was
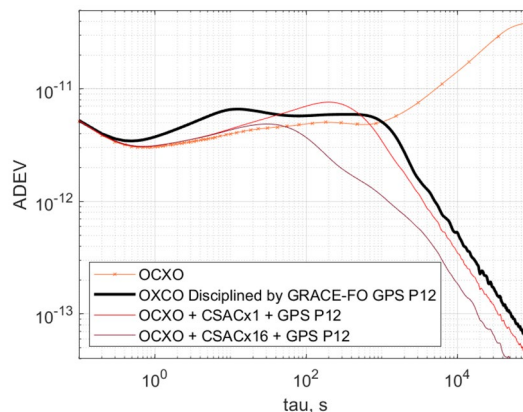

Fig. 13 ADEV for OCXO + GPS

tuned to minimize the short-term error without sacrificing too much of the long-term stability. The standard deviations and maxima are slightly worse for the OCXO systems with one and four CSACs as a result of this compromise.

It is also possible to discipline the OCXO directly from the GPS reference time using the PLL, simplifying the system and reducing cost. Figure 13 shows the result of this approach compared to the OCXO-CSAC-GPS systems. The long-term stability is not as good as the system with a single CSAC, and the short-term stability has also suffered as a result of the compromise between short- and long-term performance when selecting the PLL bandwidth. The standard deviation of the OXCO-GPS system is 2.9 ns, and the maximum error is 10.1 ns.

Conclusions

This case study provides insight into the potential performance of a heterogeneous clock system for small LEO satellites. It provides practical methods for combining the clocks to achieve optimal performance and provides a hardware architecture and methods to create low-power systems.

Table 5 shows the power for the various heterogeneous systems, along with the time error standard deviations and maxima. The power for the different systems varies from 2.4 W for the single CSAC system with the EMXO, up to 5.77 W for the system with 16 CSACs and the OCXO, demonstrating that heterogeneous systems can be created which use only a few Watts. The 16 CSACs and OCXO of the largest system occupy less than 50% of a 1U volume, based on the dimensions given in the datasheets. A single CSAC and EMXO occupy a much lower volume.

The time errors are almost entirely dependent on the number of CSACs used, with the sixteen CSAC systems approaching 1 ns of error standard deviation, and the single CSAC systems closer to 2 ns. This is because the CSAC-GPS

Table 5 Total power and time errors for various heterogeneous systems

	EMXO			OCXO			
	1 CSAC	4 CSACs	16 CSACs	OCXO Only	1 CSAC	4 CSACs	16 CSACs
FPGA	0.25	0.25	0.25	0.25	0.25	0.25	0.25
CSACs	0.12	0.48	1.92	0.00	0.12	0.48	1.92
XO	0.25	0.25	0.25	1.80	1.80	1.80	1.80
GPS	1.80	1.80	1.80	1.80	1.80	1.80	1.80
DAC	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Total, W	2.43	2.78	4.22	3.85	3.97	4.33	5.77
Time error std. dev., ns	1.90	1.50	1.10	2.90	2.10	1.50	1.10
Max time error, ns	10.50	8.70	5.10	10.10	11.00	8.90	5.10
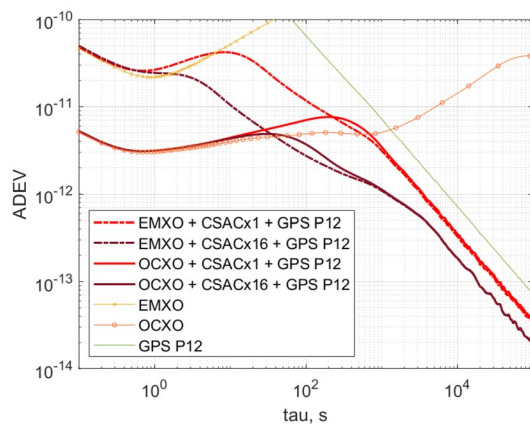
Fig. 14 ADEVs comparing OCXO + CSACs + GPS to EMXO + CSACs + GPS

combination controls the long-timeframe (> 1000 s) performance of the system, which represents the largest source of error. The use of the OCXO instead of the EMXO has almost no bearing on the time error standard deviation, but does provide significant improvement for timeframes shorter than 200 s. Figure 14 compares the ADEVs of the EMXO-CSAC-GPS systems with the OCXO-CSAC-GPS systems.

This investigation has demonstrated the feasibility of a low SWAP heterogeneous clock system that is suitable for use on small satellites. Future work is planned to implement the system in hardware and test the performance.

References

Alminde L, Laursen KK (2009) A strategic approach to developing space capabilities using Cubesat technology. In: 2009 4th international conference on recent advances in space technologies. IEEE, pp 43–47. https://doi.org/10.1109/RAST.2009.5158239

Bagala T, Fibich A, Kubinec P, Stofanik V (2016a) Enhanced clock based on the Chip Scale Atomic Clock and Dual-Mode Crystal Oscillator. In: 2016 26th international conference radioelektronika. IEEE, pp 199–202. https://doi.org/10.1109/RADIOELEK.2016.7477391

Bagala T, Fibich A, Kubinec P, Stofanik V (2016b) Improvement of short-term frequency stability of the Chip Scale Atomic Clock. In: 2016 IEEE international frequency control symposium (IFCS). IEEE, pp 1–4. https://doi.org/10.1109/FCS.2016.7546746

Best RE (2007) Phase locked loops: design, simulation, and applications. McGraw-Hill Professional, New York

Brown RG, Hwang PY (2012) Introduction to random signals and applied Kalman filtering: with MATLAB exercises, vol 4. Wiley, New York

Cash P et al (2018) Microsemi chip scale atomic clock (CSAC) technical status, applications, and future plans. In: 2018 European frequency and time forum (EFTF). IEEE, pp 65–71. https://doi.org/10.1109/EFTF.2018.8408999

Chang PC, Peng HM, Lin SY (2004) Allan variance estimated by phase noise measurements. National Standard Time And Frequency Lab Taoyuan (TAIWAN), Taoyuan

Davis JA, Greenhall CA, Stacey PW (2005) A Kalman filter clock algorithm for use in the presence of flicker frequency modulation noise. Metrologia 42(1):1. https://doi.org/10.1088/0026-1394/42/1/001

Fernández E, Calero D, Parés ME (2017) CSAC characterization and its impact on GNSS clock augmentation performance. Sensors 17(2):370. https://doi.org/10.3390/s17020370

Galleani L (2008) A tutorial on the two-state model of the atomic clock noise. Metrologia 45(6):S175. https://doi.org/10.1088/0026-1394/45/6/s23

Galleani L, Tavella P (2003) On the use of the Kalman filter in timescales. Metrologia 40(3):S326. https://doi.org/10.1088/0026-1394/40/3/312

Ge H, Li B, Ge M, Zang N, Nie L, Shen Y, Schuh H (2018) Initial assessment of precise point positioning with LEO enhanced global navigation satellite systems (LeGNSS). Remote Sens 10(7):984. https://doi.org/10.3390/rs10070984

Greenhall CA (2012) A review of reduced Kalman filters for clock ensembles. IEEE Trans Ultrason Ferroelectr Freq Control 59(3):491–496. https://doi.org/10.1109/TUFFC.2012.2219

Imken T, Castillo-Rogez J, He Y, Baker J, Marinan A (2017) CubeSat flight system development for enabling deep space science.

In: 2017 IEEE aerospace conference. IEEE, pp 1–14. https://doi.org/10.1109/AERO.2017.7943885

Kornfeld RP, Arnold BW, Gross MA, Dahya NT, Klipstein WM, Gath PF, Bettadpur S (2019) GRACE-FO: the gravity recovery and climate experiment follow-on mission. J Spacecr Rocket 56(3):931–951. https://doi.org/10.2514/1.A34326

Lombardi MA (2008) The use of GPS disciplined oscillators as primary frequency standards for calibration and metrology laboratories. NCSLI Meas 3(3):56–65. https://doi.org/10.1080/19315775.2008.11721437

Lutwak R (2011) The SA.45s chip-scale atomic clock—early production statistics. In: Proceedings of the 43rd annual precise time and time interval systems and applications meeting, Long Beach, California, pp 207–220

MicroChip (2019) Low-Noise Chip-Scale Atomic Clock Datasheet. https://www.microsemi.com/document-portal/doc_view/133866-low-noise-csac-datasheet. Accessed 5 Oct 2020

MicroSemi (2019) Quantum SA.45s CSAC, Chip Scale Atomic Clock Datasheet. https://www.microsemi.com/document-portal/doc_download/1243238-space-csac-datasheet. Accessed 5 Oct 2020

Montenbruck O, Kroes R (2003) In-flight performance analysis of the CHAMP BlackJack GPS receiver. GPS Solut 7(2):74–86. https://doi.org/10.1007/s10291-003-0055-5

NASA (2021) Physical Oceanography Distributed Active Archive Center Website. https://podaac.jpl.nasa.gov/dataaccess. Accessed 5 Oct 2020

Norton CD, Millar PS, Bauer R, Komar GJ (2016) New capabilities for Earth Science measurements with 6U CubeSat technologies. In: 2016 IEEE international geoscience and remote sensing symposium (IGARSS). IEEE, pp 5550–5552. https://doi.org/10.1109/IGARSS.2016.7730448

NovAtel. OEM719 Multi-Frequency GNSS Receiver Datasheet. www.novatel.com/assets/Documents/Papers/OEM719-Product-Sheet.pdf

Rajan RT, Bentum M, Boonstra AJ (2013) Synchronization for space based ultra low frequency interferometry. In: 2013 IEEE aerospace conference. IEEE, pp 1–8. https://doi.org/10.1109/AERO.2013.6496931

Rajan RT et al (2016) Space-based aperture array for ultra-long wavelength radio astronomy. Exp Astron 41(1–2):271–306. https://doi.org/10.1007/s10686-015-9486-6

Ramian F (2015) Time domain oscillator stability measurement Allan variance. Application note. Rohde & Schwarz, Munich

Reid TG, Neish AM, Walter T, Enge PK (2018) Broadband LEO constellations for navigation. Navig J Inst Navig 65(2):205–220. https://doi.org/10.1002/navi.234

Rybak M, Axelrad P, Suebert J (2018) Investigation of CSAC driven one-way ranging performance for CubeSat navigation. In: 32nd annual AIAA/USU conference on small satellites, 2018, Session 10 Ground Systems: 318. https://digitalcommons.usu.edu/smallsat/2018/all2018/317/. Accessed 5 Oct 2020

Segret B, Hestroffer D, Quinsac G, Agnan M, Vannitsen J, Mosser B (2018) In-flight orbit determination for a deep space CubeSat. In: 2018 IEEE aerospace conference. IEEE, pp 1–12. https://doi.org/10.1109/AERO.2018.8396756

Senior K, Coleman MJ (2014) Generation of ensemble timescales for clocks at the naval research laboratory. In: Proceedings of the 46th annual precise time and time interval systems and applications meeting, December, pp 98–105

Shah SAA, Arshad U (2013) General system design of Cubesat in LEO for IR imaging. In: 2013 international conference on aerospace science & engineering (ICASE). IEEE, pp 1–5. https://doi.org/10.1109/ICASE.2013.6785552

Stein SR (2003) Time scales demystified. In: IEEE international frequency control symposium and PDA exhibition jointly with the 17th European frequency and time forum, 2003. Proceedings of the 2003. IEEE, pp 223–227. https://doi.org/10.1109/FREQ.2003.1275093

Tkacenko A et al (2019) InSight/MarCO opportunistic multiple spacecraft per antenna (OMSPA) demonstration. In: 2019 IEEE aerospace conference. IEEE, pp 1–18. https://doi.org/10.1109/AERO.2019.8741775

Van Buren D, Palo S, Axelrad P (2019) Simulation of a high stability reference clock for small satellites with modeled GPS timing errors. In: Proceedings of the 33rd annual AIAA/USU conference on small satellites, Logan, Utah, 2019. https://digitalcommons.usu.edu/smallsat/2019/all2019/159/. Accessed 5 Oct 2020

Weaver GL, Miranian M, Garstecki JF (2012) Composite USO/CSAC timekeeping system. In: 2012 IEEE aerospace conference. IEEE, pp 1–9. https://doi.org/10.1109/AERO.2012.6187111

Wen HY et al (2019) Gravity recovery and climate experiment follow-on (GRACE-FO) level-1 data product user handbook

Weston S et al (2020) State of the art: small spacecraft technology, NASA Report

Damon Van Buren is a recent Ph.D. graduate in Aerospace Engineering Sciences at the University of Colorado Boulder. He received B.S. and M.S. degrees in Electrical Engineering from the University of Wyoming in 1991 and 1994. From 1995 to 2015, he worked in the aerospace industry developing signal processing systems for satellite ground stations and onboard processors. His industry experience includes leading the DSP developments for the Iridium NEXT onboard processor, sophisticated channelizer payloads, image compression IP, and reconfigurable computing hardware. He is president of Space Signal Sciences, LLC, which specializes in signal processing developments for space systems.

Penina Axelrad is a Professor of Aerospace Engineering Sciences and a member of the Colorado Center for Astrodynamics Research at the University of Colorado Boulder. She has been involved in GPS related research and active in the Institute of Navigation since 1986. Dr. Axelrad is a Fellow of the ION and the AIAA, and a member of the National Academy of Engineering.

Scott Palo is a Professor of Aerospace Engineering Sciences, University of Colorado, Boulder. His research interests include the design, construction, deployment, and operation of HF/VHF meteor radar systems, statistical detection, estimation, and direction of arrival algorithms, atmospheric remote sensing, global nonlinear atmospheric dynamics, and sampling/aliasing issues related to asynoptic satellite sampling.