Frequency Comparisons via GPS Carrier-phase: Jump Processing, Temperature Compensation and Zero/Short-baseline Noise-floors

Daphna G. Enzer, David W. Murphy and William A. Diener,
Jet Propulsion Laboratory, California Institute of Technology, CA, USA

BIOGRAPHIES

Daphna G. Enzer received a Ph.D. in atomic physics from Harvard University in 1996. She is currently a member of the Frequency and Timing Advanced Instrument Development Group at the Jet Propulsion Laboratory where she has worked since 2001. Her interests include development and analysis of precision measurement instruments, with past contributions to cesium-fountain space-clock prototype development; Gravity Recovery and Interior Laboratory (GRAIL); and in-house frequency stability measurement-system development. She currently works on the Deep Space Atomic Clock (DSAC) Mission.

David W. Murphy received a B.A. degree (Hons.) in physics from the University of Cambridge, in 1983, and the Ph.D. degree in physics from The University of Manchester in 1988. He is the supervisor of the Near Earth Tracking Applications Group at the Jet Propulsion Laboratory, and also works on the Deep Space Atomic Clock (DSAC) mission.

William A. Diener received a A.A. degree in aircraft maintenance from the Glendale College, CA. He joined the Frequency Standards Laboratory at the Jet Propulsion Laboratory in 1982. He currently leads testing there, and is responsible for hydrogen-maser reference standards oversight and calibration, both at the Jet Propulsion Laboratory and in the National Aeronautics and Space Administration Deep Space Network.

ABSTRACT

GPS carrier-phase receivers have been used to produce high quality time/frequency comparisons between clocks and frequency standards. Current ways of processing the carrier-phase data, however, result in phase jumps and data gaps that make the data difficult to use effectively and repeatedly for long datasets. In addition, temperature fluctuations generally perturb receiver output detrimentally. We describe the steps of a simple but effective automated algorithm for comparing frequencies that can compensate for data gaps, day boundary jumps, and other phase jumps as well as for temperature disturbances. We describe process features and focus on zero-baseline noise-floors, measuring Allan deviations of the Deep Space Atomic Clock mission's receiver down to 1×10^{-17} at 5×10^{5} seconds, with an upper confidence level of 3×10^{-17}. The zero-baseline noise floors we present give confidence in our algorithm and temperature compensation's robustness.

INTRODUCTION

The atomic clock community is often faced with the need to compare outputs from different frequency standards on a long-term basis. Measurement systems must be robust for months or even years and have noise floors sufficiently low not to limit the comparison at hand. At the Jet Propulsion Laboratory (JPL), we have typically accomplished this type of comparison between local standards with specially developed phase/frequency measurement systems [1], [2]. Meanwhile, longer baseline stability measurements (e.g. between local clocks and UTC(k) laboratories) have traditionally been accomplished in the community via traveling clocks, two-way satellite time and frequency transfer (TWSTFT), GPS code time transfer, or more recently via long-distance fibers or GPS carrier-phase (plus code if needed for time accuracy) [3]-[10].

At JPL, we have long used GPS code time transfer to measure and calibrate atomic frequency-standards and clocks at Deep Space Network [11] sites in the United States, Spain, and Australia. We have also performed carrier-phase based comparisons, in particular, between one of JPL's ultra-stable mercury ion-clocks and UTC(k) laboratories [12]. These comparisons, however, require manual data processing and boundary/phase jump corrections. Our new post-processing algorithm described herein expands upon this work in preparation for stability characterization needs of the upcoming Deep Space Atomic Clock (DSAC) mission, which aims to demonstrate mercury ion-clock operation in space [13], [14]. New achievements include 1) improved processing and automation, 2) successful implementation of temperature calibration and compensation for GPS receivers in

varying thermal environments, and 3) characterization of the GPS frequency-transfer noise floor over zero and short measurement baselines.

Carrier-phase receivers offer the opportunity for precision frequency measurements, but unfortunately, they do not easily provide the continuous set of phase versus time data typical of benchtop (local) measurement systems such as in [1] and [2]. Much processing has to occur, and state-of-the-art first-stage processing still leaves day boundary and batch-boundary jumps, plus perhaps additional jumps and/or data gaps due to resets or other issues with the extracted clock solutions [3]-[5]. We present here an overview and basic validation of our 5-stage algorithm to account for these commonly encountered jumps, using interpolation. We focus on the conceptual framework, the achievable zero-baseline noise floors, and on the reproducibility of overlaid results from one dataset to the next. The thrust of these measurements was to show the robustness of our automated processing and to measure the noise-floor of our receivers, so we present this as a study of zero/short-baseline common-clock noise-floors. Carrier-phase receivers are commonly used to compare clocks across great distances, but we focus, as a first step, on validating that our simple automated algorithm is not degrading the receiver zero-baseline noise-floor. As a convenient byproduct, we found these short-baseline noise floors low enough to enable clock comparisons across the JPL campus with this technique, when stable reference clocks were not available in remote buildings (but state-of-the-art GPS roof antennae were).

We implement our algorithm after an initial processing stage with the JPL-developed GIPSY-OASIS (GPS-Inferred Positioning System and Orbit Analysis Simulation Software) [15] software in which the offset $x(t)$ between the receiver clock and either a remote or local reference clock has already been determined. Note however, that our technique should be easily extendable for use with GPS (and/or GNSS) processing software other than GIPSY-OASIS (referred to in this paper simply as GIPSY). Note also, that we used GIPSY solutions with robust integer ambiguity resolution via the wide-lane phase bias technique in [15].

Our procedure detects data gaps and delay jumps outside of the measured noise envelope, adjusts the $x(t)$ delays across each jump, and extracts a valid Allan deviation [16] for available data. We do not claim to measure anything about the frequency source during times of data outage, but for the data that are present, a rigorous Allan deviation is computed with the main caution being that our jump detection could mistakenly remove a true jump in the source signal. However, most traditional frequency measurement systems have similar trouble tracking large jumps, and often such large jumps are not physically plausible for the clock under test. We plan to study the filtering criteria for each clock under test and evaluate the likelihood of any of the filtered jumps being physically plausible. Thermal compensation is optionally included during one of the stages of the automated processing. We describe this compensation technique, validation data, and calibration of the DSAC receivers' sensitivities.

The next section covers our 5 stage processing technique, leaving details of the temperature compensation to be covered in the third section. The third section presents our measurements and starts with an overview of the receivers used, followed by an ideal environment zero-baseline noise floor, temperature impact and compensation, an evaluation of the different receivers' performances, short-baseline effects, and a brief discussion about expanding this validation to long-baseline data.

5 STAGE PROCESSING TECHNIQUE

Figure 1 shows an example report from our new processing algorithm showing the five stages. Each row shows time-delay ($x(t)$), fractional frequency ($y(t)$) and overlapping Allan deviation (OAD) for one of the five stages.

In the first stage, our automated data fetching program creates an $x(t)$ time series over the time range requested (typically spanning multiple days) from daily $x(t)$ data generated by other processing software that runs GIPSY. Our GIPSY processing can process two types of delay data: single receiver data and receiver pair data.

Single receiver data is the receiver-under-test's delay relative to a reference clock, and is determined using GIPSY in its single receiver daily static Precise Point Positioning (PPP) mode in which the receiver-under-test's position is determined once a day using the ionosphere-free pseudo-range and carrier-phase observables, sometimes called PC and LC data. This analysis uses JPL's GPS orbit and clock products in which the GPS satellite transmitter clocks are determined relative to a ground-based reference receiver chosen during the creation of these products. This approach allows us to determine the clock offsets $x(t)$ between the receiver-under-test and the reference clock, typically, and in this paper, at a 300s cadence, in what we call single receiver delay data, because there is only one receiver under test. Fixed receiver delays are ignored because we only use this technique for frequency measurements. The reference is chosen (for all GIPSY users that day) from a list of reference clocks

steered to UTC. During the time period covered by this paper the reference was usually a United States Naval Observatory (USNO) clock feeding either the USN7 or AMC2 receiver.

Receiver pair data, on the other hand, measures the relative delay between two receivers under test, and we derive it directly from GIPSY by assigning one of the receivers to be the reference clock when performing zero/short-baseline processing of the GPS data. We follow an approach that is typically used to generate high accuracy relative positioning between GPS antennae that have very short baselines, taking advantage of common-view zero-baseline cancelation of various effects [17]. For this processing, the GPS satellite orbits are fixed to their values in JPL's GPS orbit and clock products, but the transmitter clocks are solved along with the ground-based receiver that is not the reference clock, typically, and in this paper, at a 30s cadence. By virtue of the zero/short baseline, this method removes the common-mode troposphere and ionosphere delay along each transmitter-receiver line of sight since they get absorbed into the estimated transmitter clock, and so allows for a better estimate of the clock difference between the two ground receivers. Some of this clock difference will be due to different internal delays in the receivers-under-test, even if the same external clock feeds both, but since we are only interested in relative frequency measurements, this is immaterial for stationary delays. Furthermore, since the frequency-dependent ionospheric delay is accounted for in the estimated transmitter clocks, a single carrier-phase frequency such as L1 can be used instead of the LC/PC ionospheric-free combination. Doing so further reduces noise in the frequency comparison for what we call pair data in this paper.

Row 1 of Fig. 1 shows this stage of processing zero-baseline pair data, with one large jump in $x(t)$ after a data gap. Delay jumps are common after data gaps if the gap was due to loss of lock with GPS satellites. Time along the x-axis is labeled in days such that the values correspond to the date, at least until a month crossover occurs. Since this run started in December, we see it covers the timespan Dec. 21 through Jan. 2 (end-of-day). The fractional frequency plot is merely a point-by-point derivative of the delay plot, and the Allan deviation calculation takes into account data gaps with the following method: While most Allan deviation calculations assume that the $x(t)$ data are equally spaced in time, this is violated when there are data gaps. The Allan deviation that we use pads the $x(t)$ data with placeholder data in order to get equally spaced $x(t)$ data, solely to bookkeep the gaps for the Allan deviation calculation. This placeholder data is not used for any actual calculations. It is merely a convenient way to bookkeep missing data points and simplify the compilation of the $x(t-\tau)+x(t+\tau) - 2*x(t)$ second difference terms. All these terms are computed whether they have real or placeholder data, but the final sum that goes into the Allan deviation only adds the terms containing real data. In fact, the placeholder delays could be assigned any errant value, because they do not contribute to the Allan deviation results. Of course, this Allan deviation does not reflect all events that happened during the timespan, because no information is available during data gaps, but it does represent the true noise character during the times that we have data.

The second stage of processing removes delay jumps associated with data gaps, taking into account the slope in $x(t)$ (assumed to be median $y(t)$). The day-boundary jumps are now apparent (Row 2). Temperature compensation is optionally performed before jump correction in this stage, although for this particular case, it does not have a visible impact because the receivers were located in a thermally controlled environment. Again, the Allan deviation continues to account for missing data by not including those second differences in the sum, as is the case for all Allan deviation calculations in the 5-Stage processing. (Our temperature compensation will be described in the next section in conjunction with the calibration and compensation data we present.)

The third stage of processing removes the median $y(t)$ which is insensitive to $x(t)$ outliers as opposed to the best-fit $x(t)$ slope which could become very corrupted by even a single $x(t)$ outlier. This frequency removal is tracked and reported, along with a later frequency removal in order to maintain connection to an absolute frequency readout. This stage also shows a red 'x' for the $y(t)$ points that are flagged for jump correction in the next stage of processing. Since no actual jump-correction is performed in this step, the day boundary jumps remain dominant for this particular data set. If the two receivers under test had different source frequencies, this could be the first stage to reveal the day-boundary jumps.

The fourth stage of processing corrects for the jumps identified in Stage-3 by inter-quartile range (IQR) filtering. All $y(t)$ with $|y(t)\text{-median}| > \text{iqr} * \text{IQRF}$ are filtered out, where median = median($y(t)$), iqr = $y(t)$ inter-quartile range, and IQRF = IQR Factor chosen by the user. The median $y(t)$ should be zero at this stage since it was already removed in Stage-3, which is a necessary precursor to Stage-4. This filtering process is much less sensitive to outliers than filtering based on the standard deviation would be (since the standard deviation is heavily dependent on the very outlier data we want to remove). This jump correction step eliminates day boundary jumps as well as any other large enough delay glitches (other than those caused by data gaps). This is the place that the user needs to be careful not to filter out events that are actually caused by the frequency source under test. The goal here is to set the filtering identification strong enough to remove all GPS glitches, but weak enough to let through all

jumps in the clock being tested. Detailed knowledge of the source/clock in use and understanding the possibility/likelihood of certain sized jumps helps the user determine this jump threshold. The user chosen IQR Factor is shown on the Stage-4 plot title, and experience has shown that 4.0 and 10.0 are good IQRF choices for typical single receiver and pair data, respectively. For the noise-floor data shown in Fig. 1, the jump threshold was 3.7×10^{-13}, which would be quite a low value to use for actual clock data (*i.e.* could filter out real clock jumps) However, clock data would have a higher level of noise, so the threshold would scale accordingly (for a constant IQR Factor). In a future upgrade to this processing technique, we plan to separate the day boundary jumps from the non-day-boundary jumps in order to set separate thresholds and also to set minimum *y(t)* filtering thresholds based on physical knowledge of the clock under test.

The fifth and final stage of processing removes a linear best fit for *x(t)*, which is yet another fractional frequency removal that is tracked along with the first frequency removal in Stage-3. Since this example dataset is a common-LO measurement, both receivers have near the same frequency, so there is not much of a total bias frequency to remove in these 2 stages. However, the median removal of Stage-3 could remove too much or too little *x(t)* slope due to outliers or due to small differences between median *y(t)* and best fit *x(t)* slope, so often this last stage still has a significant removal and is the first to reveal the true scatter in *x(t)*.

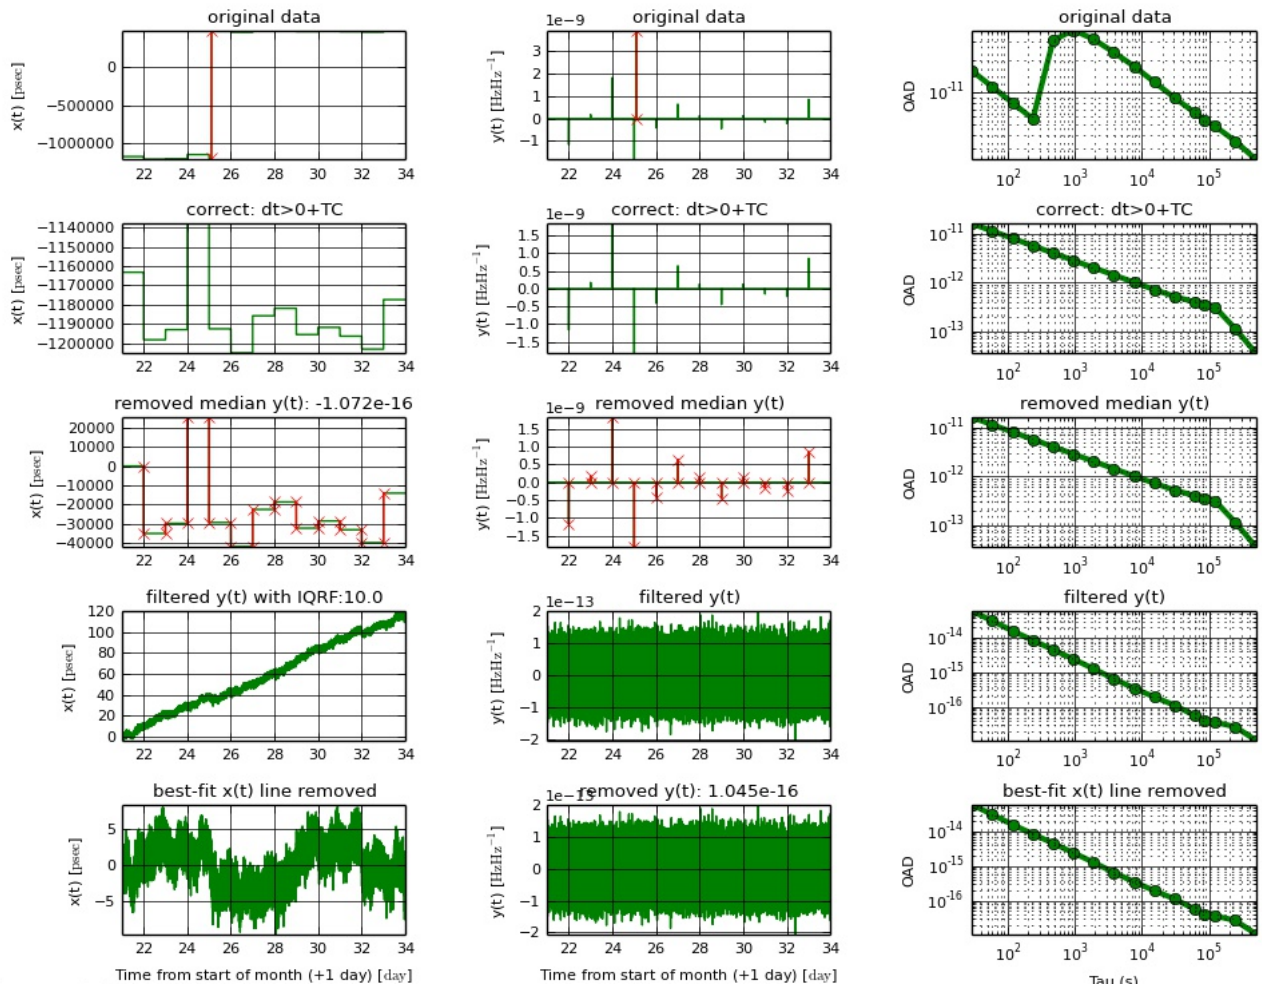

Figure 1. 5-stage processing example report, Run1 in Table II of the Appendix. Each row corresponds to one stage of processing, showing: delay (*x(t)*), fractional frequency (*y(t)*) and overlapping Allan deviation (OAD).

MEASUREMENTS

Overview

In the subsections that follow, we present noise floors for this technique using three different GPS receivers. We focus on pair noise-floors for which we compare the delay of one receiver to the delay of another, both sourced by the same LO. This is motivated by wanting to measure the performance of the DSAC flight receiver and to exercise our temperature compensation technique on our most sensitive data. Most pair results shown are zero-baseline, *i.e.* making use of the same antenna, but towards the end we show the impact of short-baseline where the two antennae were located on rooftops of different buildings at JPL. Finally, we comment on how to relate these zero/short-baseline measurements to expectations for long-baseline (single receiver) measurements.

Receivers

The three receivers we used are shown in Table I. JPLT is an Ashtech Z12T, located in our thermally controlled Frequency Standards Test Lab (FSTL) at JPL. EM and FM are both TriG-based receivers [18] built for DSAC, but referred to here as TriG-Lite because they lack the Linux-based Science Processor. EM is the engineering model used only for testing, and FM is the unit currently on the spacecraft, awaiting launch, but also used in many of the tests presented. These two DSAC receivers have similar designs and similar performance, but different thermal sensitivities. We transported them between several different buildings on the JPL campus for these measurements, as indicated in Table II of the Appendix.

TABLE I
GPS RECEIVERS TESTED

Receiver Label	Receiver Type/Model	Location
JPLT	Ashtech Z12T	JPL FSTL
EM	DSAC Engineering Model	JPL (varied)
FM	DSAC Flight Model	JPL (varied)

Ideal-environment noise-floor: zero-baseline, common LO, common antenna

The results shown throughout this paper come from noise floors measured during DSAC ground testing at JPL. These were obtained to check baselines before, in-between, or after various DSAC clock characterization campaigns and therefore, were often no longer than a few days in duration. Figure 2 shows the longest zero-baseline we performed with the DSAC receivers (FM and EM). Both were in a temperature-controlled environment, so temperature compensation was not necessary. As such, this figure represents the ideal-environment pair-measurement for these receivers: common LO, zero-baseline, little temperature variation. The noise-floor result (shown in blue) represents an upper limit to the measurement noise of each DSAC receiver. Furthermore, it validates that our algorithm can handle typical jumps and data gaps (12 day-boundary jumps and one short 420 s data gap and jump, here) and reproduces the expected result of zero, down to 1×10^{-17} at 5×10^{5} s, with an upper confidence level of 4×10^{-17}. All confidence levels in this paper are reported at the 1-sigma level and are calculated from [16] with an assumption of white frequency noise.

Even though both receivers were sourced from the same H-maser LO, they each required a 20.456 MHz input which came from a 4-channel synthesizer (referenced to the maser). One output channel went to the FM, while another went to the EM. The black dashed line in Fig. 1 shows a separately measured channel-versus-channel synthesizer noise-floor, included here to confirm that the synthesizer component of the LO had negligible impact. (This synthesizer Allan deviation was measured by offsetting one of the 20.456 MHz channels by 1 Hz and mixing the two channels together, when they were not being used as inputs to the EM and FM.)

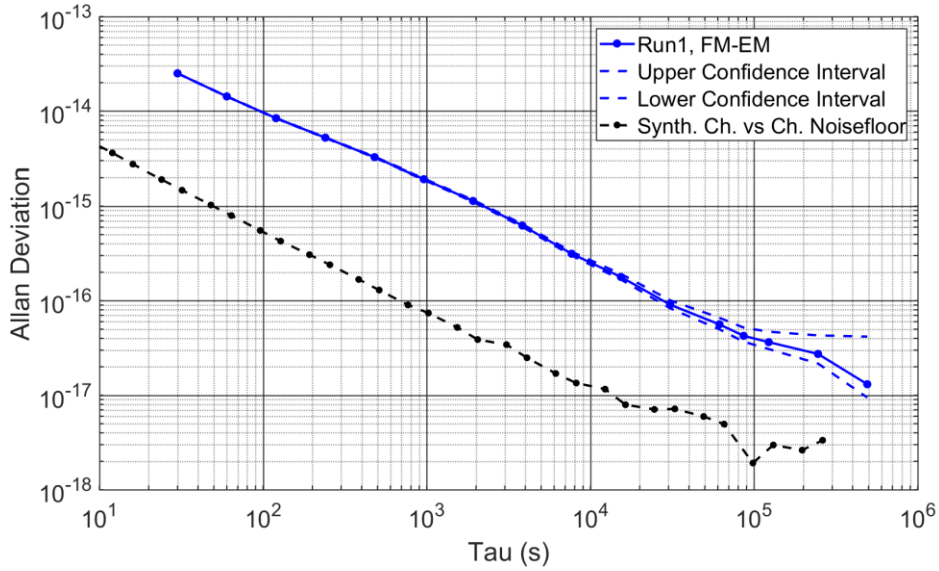

Figure 2. Allan deviation for zero-baseline, common-LO noise-floor using DSAC TriG-Lite receivers (FM and EM) in a temperature stable environment. Blue lines show FM-EM pair data while the black dashed line shows the noise-floor contribution due to the synthesizer that fed both TriG-Lite receivers. Time range and receiver locations for the FM-EM noise floor (Run1) are in Table II of the Appendix. The synthesizer Allan deviation shown is non-overlapping while all the other legend entries are overlapping.

Temperature impact and compensation

In less ideal environments, and in particular, in time-variable environments that DSAC will encounter, temperature can have a large impact on the measured output of a receiver. Figure 3 shows an example of the impact due to receiver warmup and room temperature variation, a 5 degree Celsius span in this case, and the effectiveness of our temperature compensation. It shows one of our FM-EM zero-baseline runs (Run4 of Table II), both with and without compensation. Legend entries marked with 'TC' have temperature compensation applied. Comparing the dark blue versus cyan curves, we see that compensation is necessary to achieve low Allan deviations such as shown here. (Furthermore, this Allan deviation is shown to be near that of a temperature stable run in Fig. 5.) Each 'TC' label is followed by '-EM' or '-FM' to indicate which receiver was temperature compensated. Typically, the EM receiver was on a thermally controlled plate, so compensation was not as necessary. Fig. 3 also displays concurrent single receiver FM data in order to convey that the temperature impact was almost big enough to degrade this lower performing single receiver curve as well (red versus purple curve).

Temperature compensation is (optionally) applied in Stage-2 of our post-processing algorithm, and as can be seen, dramatically improves the Allan Deviation. Since our goal is to use these receivers to measure the frequency stability of clocks, any extraneous delay or frequency variation due to other causes needs to be either minimized or calibrated out to get the best performance. Both the FM and EM TriG-Lite receivers have significant temperature-dependent internal delays that need to be calibrated out if there are significant temperature variations. For both these receivers, we undertook a series of calibration runs (see for example Fig. 4) which lasted multiple days in which the receiver under test (either the FM or EM) had its temperature varied by many tens of degrees while a reference receiver was held at a steady temperature. Both the receiver under test and the reference receiver were driven by a common clock, and $x(t)$ between the two was determined. Furthermore, the TriG-Lite receivers produce internal temperature measurements (T) as auxiliary data. The $x(t)$ and $T(t)$ data obtained during calibration runs show strong non-linear dependence on one another which can best be captured by fitting $x(t)$ to a quartic polynomial in T whose coefficients then become calibration parameters. The FM and EM have similar but different polynomial coefficients. Furthermore, the FM receiver had two calibration runs in August 2014 and November 2015, and the coefficients produced in the two runs separated by over a year were not significantly different, giving us confidence that the calibration coefficients determined on the ground can successfully be applied on orbit. Thus, when we say we temperature compensate a receiver, we use the polynomial calibration coefficients for a particular receiver along with the measured receiver temperature to determine the temperature-dependent receiver delay at every epoch which is then subtracted from the originally determined $x(t)$ to produce a temperature compensated $x(t)$.

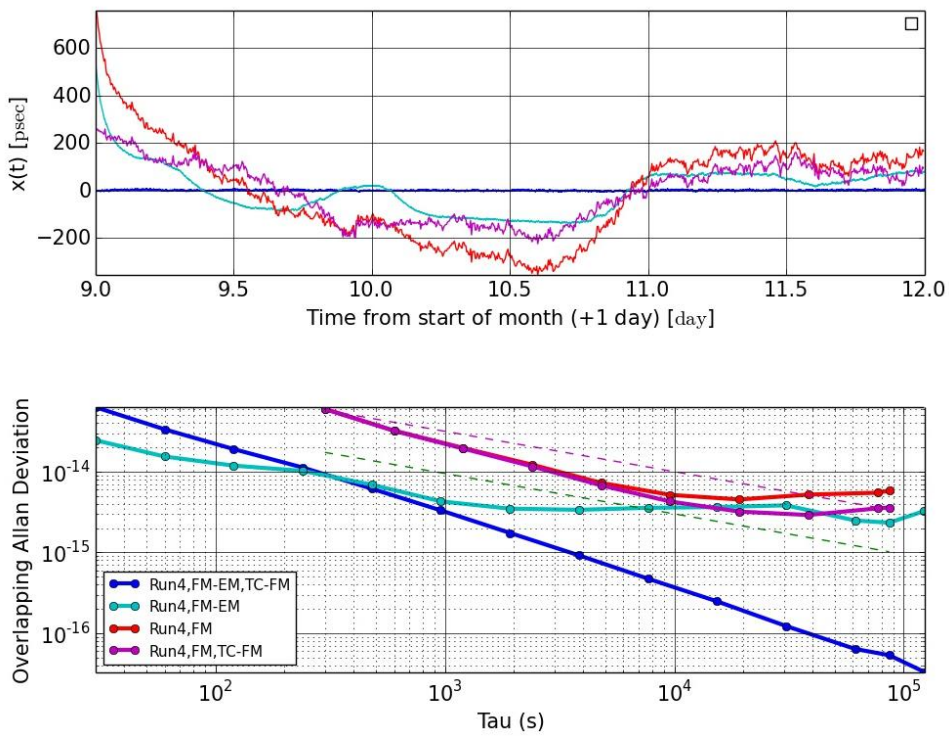

Figure 3. Delay versus time (above) and Allan deviation (below) for FM-EM and FM in the date range covered by Run4 (See Table II in Appendix). Temperature fluctuations were just 5 degrees Celsius, but still, the uncompensated FM-EM Allan deviation (cyan) was severely impacted.

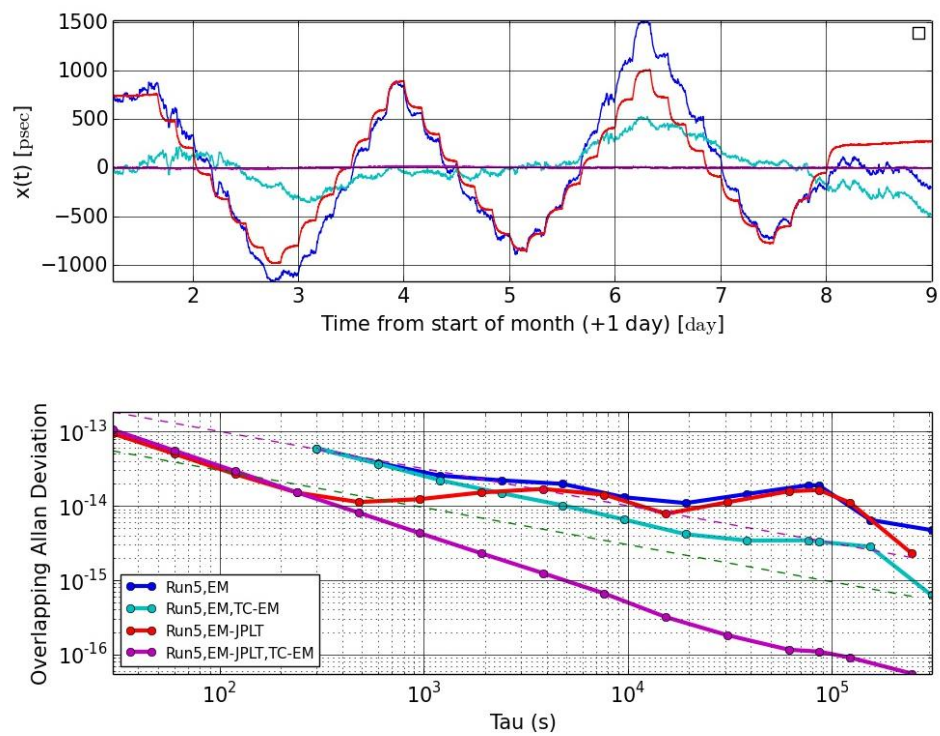

Figure 4. Delay versus time (above) and Allan deviation (below) for EM and EM-JPLT in the date range covered by Run5 (See Table II in Appendix). The EM receiver was in a thermal chamber during this time period, which explains the large step-like EM and EM-JPLT *x(t)* excursions in the top plot (red and dark blue curves). Corresponding Allan deviations in the lower

plot show a big impact due to temperature, but the compensated curves (magenta and cyan) match the expected zero- and long-baseline noise floors. EM temperatures varied 45 degrees Celsius over the course of this thermal calibration run.

Reproducibility

Figure 5 shows FM-EM zero-baseline noise floors from four different runs in order to examine their reproducibility. It overlays Runs 1 and 4 already shown in Figs. 1-3 plus two more (Runs 2 and 3). As illustrated already by Figs. 3 and 4, if temperature is not properly compensated, the Allan deviation degrades. Similarly, if day boundary and other delay jumps are not properly corrected, the Allan deviation also degrades. Thus, the reproducibility of these overlaid Allan deviations, measured in different locations at different times with different temperature variability is a validation of both the jump correction and temperature compensation at the level of agreement shown. FM receiver temperatures for these measurements were varying at most 5 degrees Celsius (due to warmup and/or room temperature), but the EM was on a thermally controlled plate.

The first two legend entries show uncompensated versus compensated Allan deviations for the temperature stable Run1 data of Fig. 2 that did not require compensation. Comparing these illustrates how our compensation procedure can degrade the Allan deviation a small amount for low tau. In most instances, however, room temperature varied too much to avoid temperature compensation. However, if the low tau degradation were really a problem for our needs, we could improve it by smoothing the temperatures before applying temperature compensation. For DSAC, other effects such as atmospheric fluctuations above the reference receiver will limit our Allan deviation on those timescales, so we were not concerned. More detail about this Run1 data set can be seen in Fig. 8, which shows the delay versus time curves for these two legend entries along with others from that run.

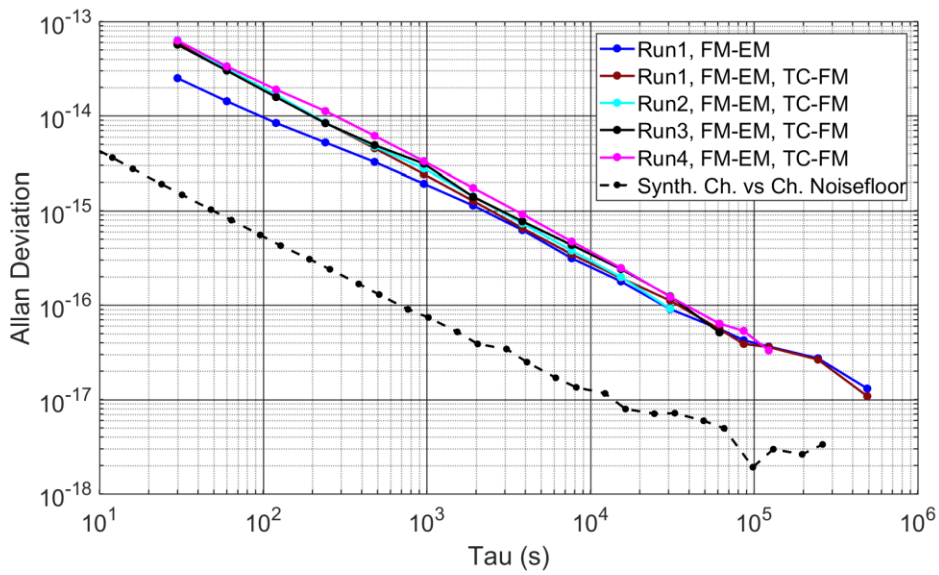

Figure 5. FM-EM Allan deviations for co-located receivers in one of three different buildings on the JPL campus. See Table II in Appendix for locations and date ranges. Dashed line is the same synthesizer channel-versus-channel noise contribution described in Fig.2 (still negligible). The synthesizer Allan deviation shown is non-overlapping while all the other legend entries are overlapping. These pair data include the following number of jumps (all day-boundary): Run1: 12, Run2: 0, Run3: 1, Run4: 2. They also include the following data gap durations (in seconds): Run1: 420, Run2: no gaps, Run3: 510, 360, and 390, Run4: 60. Finally, they also include the following FM temperature excursions (in degrees Celsius): Run1: 0.2, Run2: 0.2, Run3: 0.4, Run4: 5.

Reproducibility using other receiver pairs

Figure 6 shows another zero-baseline reproducibility plot, this time for pair measurements involving the JPLT (Ashtech) receiver. Again, these noise floors required temperature compensation as indicated to achieve near overlap, and again, receivers that were not temperature compensated were in a thermally controlled laboratory (FSTL) or on a thermally

stabilized plate or both (EM was on plate for Runs 7 and 8, but not for Runs 5 and 6.). Even with our compensation/temperature control, the flicker at 10^{-16} is likely due to residual temperature effects and/or other environmental factors. Receiver temperature excursions were at most one degree Celsius except for Run5, which was the thermal calibration run of Fig. 4.

FM-JPLT and EM-JPLT noise floors in Fig. 6 appear roughly equivalent and are limited by JPLT, *i.e.* Fig. 5 shows better performing pairs. This is seen more explicitly in the comparison plot of Fig. 7.

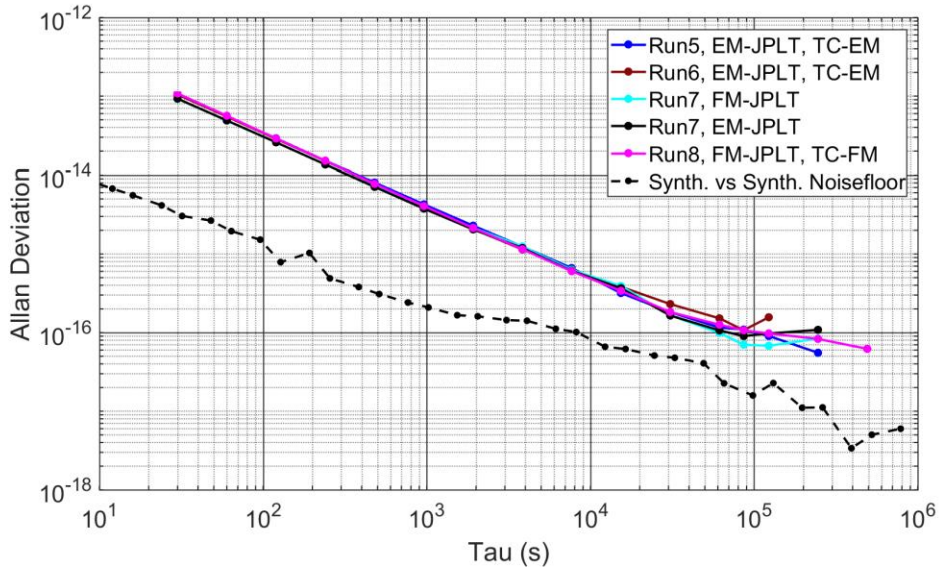

Figure 6. FM-JPLT and EM-JPLT Allan deviations for receiver pairs co-located in JPL's FSTL. See Table II in Appendix for date ranges. Dashed line is the Allan deviation for one 20.456 MHz synthesizer beat against another, 1Hz detuned and referenced to the same maser. This curve represents a worst-case synthesizer noise contribution for the noise floors shown which only made use of one channel of one synthesizer. (JPLT was driven directly by quadrupling the maser's 5 MHz and did not need these synthesizers.) The synthesizer Allan deviation shown is non-overlapping while all the other legend entries are overlapping. These pair data include the following number of jumps, all day-boundary unless indicated: Run5: 9 (2 were non-day-boundary), Run6: 2, Run7: 9 (2 were non-day-boundary), Run8: 21 (2 were non-day-boundary). They also include the following data gap durations (in seconds): Run5: 720, 60, and 960, Run6: no gaps, Run7: 420 (for FM-JPLT curve only), Run8: 390, 31110, and 60. Finally, they also include the following temperature excursions for the compensated receiver (in degrees Celsius): Run5: 45, Run6: 0.3, Run8: 1.

Comparison of receiver pairs

Figure 7 overlays FM-EM zero-baseline noise floors with FM-JPLT ones to explore the performance of the various receivers. (EM-JPLT curves were similar to FM-JPLT curves in Fig. 6, so not shown.) It is clear that the FM and EM receivers produce a slightly lower noise floor than JPLT. The lowest Allan deviation measured is on the temperature compensated FM-EM curve: 1×10^{-17} at 5×10^{5} s, with an upper confidence interval of 3×10^{-17}, however, Figs. 5 and 6 give a sense of the reproducibility of the various curves and show some curves up to 20% higher.

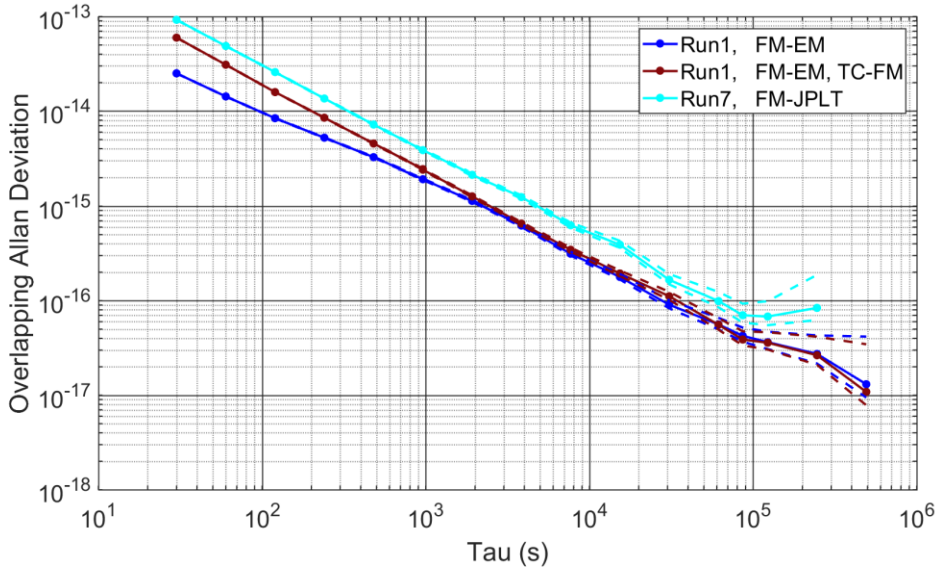

Figure 7. Comparison plot showing representative FM-EM versus FM-JPLT zero-baseline noise floors (solid lines) with 1-sigma confidence intervals (dashed lines).

Excess noise from unremoved jumps

Figure 8 shows a more complete picture of Run1 (already seen in Figs. 1, 2, 5 and 7) in order to visualize the excess noise in single receiver (long-baseline) data and to examine excess noise due to unremoved jumps. It shows *x(t)* delay versus time and Allan deviation data for various receiver pairs as well as for single receivers, all collected simultaneously and color coordinated between the top and bottom plots. The single receiver Allan deviations (first four legend entries) are dominated by long-baseline GPS time-transfer noise, and perhaps a bit of maser LO drift. Therefore, they all show the same noise floor, at or below the $10^{-12}/\sqrt{\tau}$ guideline drawn in. The FM–EM curves already shown in Figs. 1, 2, 5, and 7 are replotted here (two lowest Allan deviations) to be able to examine their corresponding delay versus time curves above. The last receiver pair shown, FM–JPLT (black), has slightly worse Allan deviation than those shown in Fig. 6 because of the jumps visible in the delay versus time plot above. The delay curves visually illustrate the extra noise in the (long-baseline) single receiver plots as compared to the (zero-baseline) pair plots, as well as of course, the delay jumps in the FM-JPLT data.

We do not usually see such delay jumps, but Run1 took place during a time in which several of the receiver power supplies and antenna signal distributions may have changed, so we were not surprised that a few glitches occurred, and we did not focus on tracking down their sources. We include these data for two reasons. First, including them shows the level of Allan deviation degradation we get from these types of small residual jumps, seen more explicitly in Fig. 9 which plots the FM-JPLT and EM-JPLT curves of this dataset (Run1) along with the best of Fig. 6 (Run7) plus a 4-day subset of Run1 that was chosen to avoid these glitches (Run10). As suspected, after removing the glitches, Run10 is as good as or better than the best of Fig. 6 (Run7). Second, we note that we do not see this type of degradation on the FM-EM Allan deviations from the same time range (Run1). If we did, the Run1 curve in Fig. 5 would be degraded above the level of the others. We deduce that either the FM and EM both had jumps that subtracted away very well in processing, or just JPLT had the jumps. Whichever it is, the resulting FM-EM Allan deviations shown in Fig. 2 are still representative of typical noise floors in the presence of potential glitches.

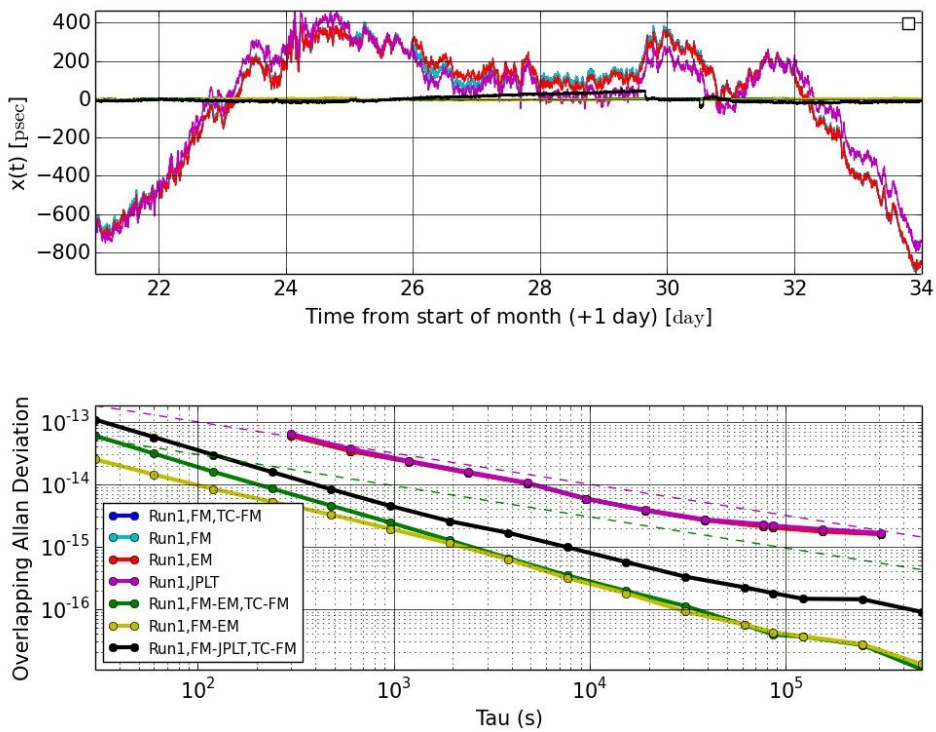

Figure 8. Delay versus time (above) and Allan deviation plots (below) for various single receivers and receiver pairs in the date range covered by Run1 (See Table II in Appendix). The single receiver delay and Allan deviations are measured against the standard reference clock used by GIPSY, which changed between three different clocks over this time range: USN7, AMC2, and WTZR. (Transitions between reference clocks are done at the day boundary, and any resulting delay jumps are removed by our processing.) Guidelines (dashed) are drawn at $10^{-12}/\sqrt{\tau}$ and $3x10^{-13}/\sqrt{\tau}$.

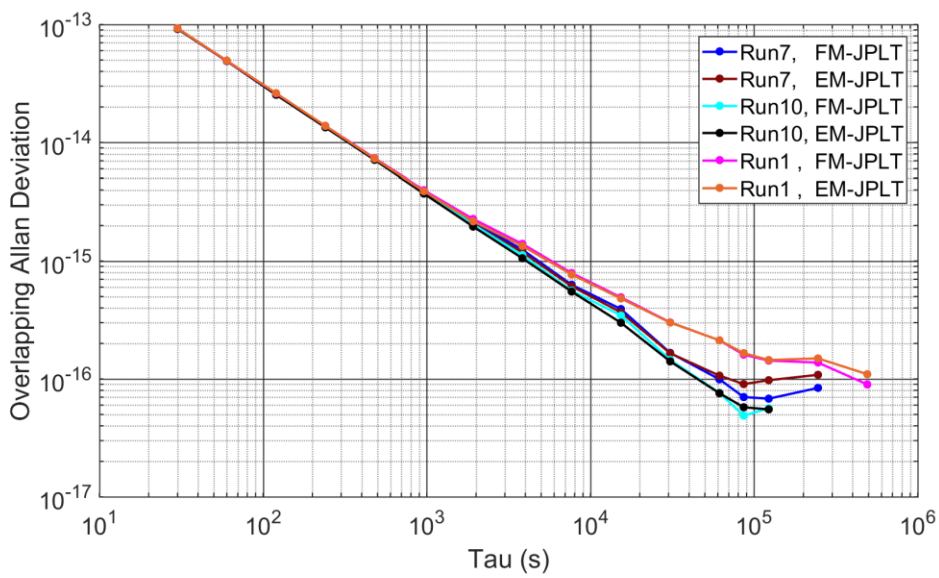

Figure 9. Best FM-JPLT and EM-JPLT Allan deviations from Fig. 6 (Run7) versus a 4-day sub-span picked to eliminate glitches (Run10). Run1, which incorporates all the dates of Run7 and Run10 and more, is also displayed. Eliminating the glitches improves the Allan deviation.

Zero- versus short-baseline

Figure 10 shows the impact of separating the two receivers in a pair across different building on the JPL campus. For the three short-baseline measurements displayed (brown, cyan, and black solid lines), the FM receiver was in either Bldg. 1 or Bldg. 2 (see Table II) while JPLT remained in FSTL. These FM-JPLT noise floors all reveal extra short-baseline noise over the zero-baseline Allan deviations of Fig. 6. However, for these measurements, the maser LO was sent to the remote building via already existing telecom fibers (*i.e.* not via a dedicated quiet link). Round trip link noise curves for both buildings are shown (dashed lines), but they could be over-estimates of the one-way links since they are round-trip links. However, since 1) they were not measured concurrently with the FM-JPLT short-baselines 2) temperature variations those days could have made their noise contributions still higher, and 3) they do approach their corresponding short-baseline curves at a few times tau, we cannot rule out that some link noise may have contributed to the short-baselines reported. In any case, the true short-baseline FM-JPLT noise floors are *no worse* than what is shown here, which is still much better than the single-receiver noise floors that include long-baseline GPS time-transfer noise. These carrier-phase comparisons have proven to be a useful way to measure clocks across the JPL campus in buildings that do not have stable references (but do have state-of-the-art carrier-phase antennae on their roofs). Note we did not need to address antenna or cable temperature sensitivity to achieve the noise-floors reported here, so any differential sensitivities were not big enough to impact the Allan deviations shown. Weather station plots from nearby reveal 5 to 15 degree Celsius diurnals in the time periods covered.

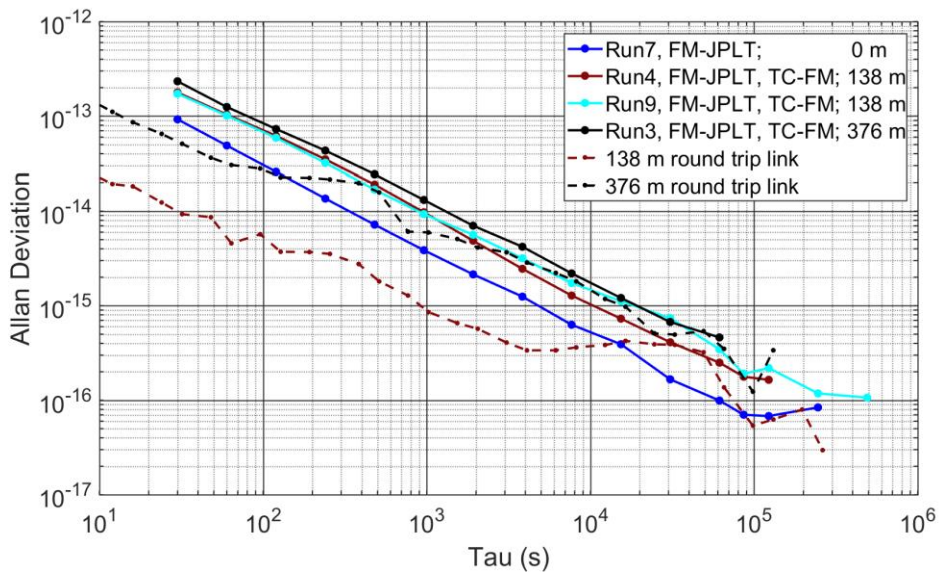

Figure 10. Comparison plot of FM-JPLT over three baselines: Figure 6's best zero-baseline (dark blue), two 138-m short-baselines (brown and cyan) and a 376-m short-baseline (black). Zero-baseline has the lowest noise floor, and the three short-baselines are not too different from each other. Dashed lines show the round trip link noise for each of the buildings. Link noise curves are non-overlapping, while all others shown are overlapping.

Long-baseline single receiver data

In Fig. 11, we overlay single receiver (long-baseline) Allan deviations covering all the timespans investigated in this paper. They show less reproducibility than the pair data overlay plots, but this could be due to several causes. We did not set out to measure the ultimate noise floor of the single receiver data because that would have required a long-term stable clock better than the expected noise-floor. Since we used a H-maser as our LO, we cannot rule out its contribution to these Allan deviation variations. More of concern though, it is possible that our jump correction is less effective for these higher noise datasets since higher *x(t)* noise means higher random walk introduced at each day boundary and gap correction is potentially degraded by less accurate knowledge of the median *y(t)*. To our advantage, however, these higher noise datasets are more forgiving to leaving residual glitches in place, because again, the higher noise one day Allan deviations (one day is the focus for DSAC) are not as sensitive to small jumps in the 10^{-17} range as are the lowest noise-floor pair Allan deviations. Also, most importantly for DSAC purposes, the fact that all the Allan deviations in Fig. 11 lie at or below the $10^{-12}/\sqrt{\tau}$ guideline out to a day gives confidence that the algorithm and temperature correction will not degrade the measured Allan deviations beyond this level. (This $10^{-12}/\sqrt{\tau}$ guideline was the assumed long-baseline time transfer noise level that went into DSAC planning.) Finally, as evidenced by the single receiver curves in Figs. 3 and 4, temperature excursions that pushed these Allan

deviations above the $10^{-12}/\sqrt{\tau}$ guideline, did come down below the line with our temperature compensation. Furthermore, since on-orbit, the DSAC receiver will be subject to less of the un-calibratable tropospheric fluctuations, we expect better single receiver performance compared to testing it on the ground.

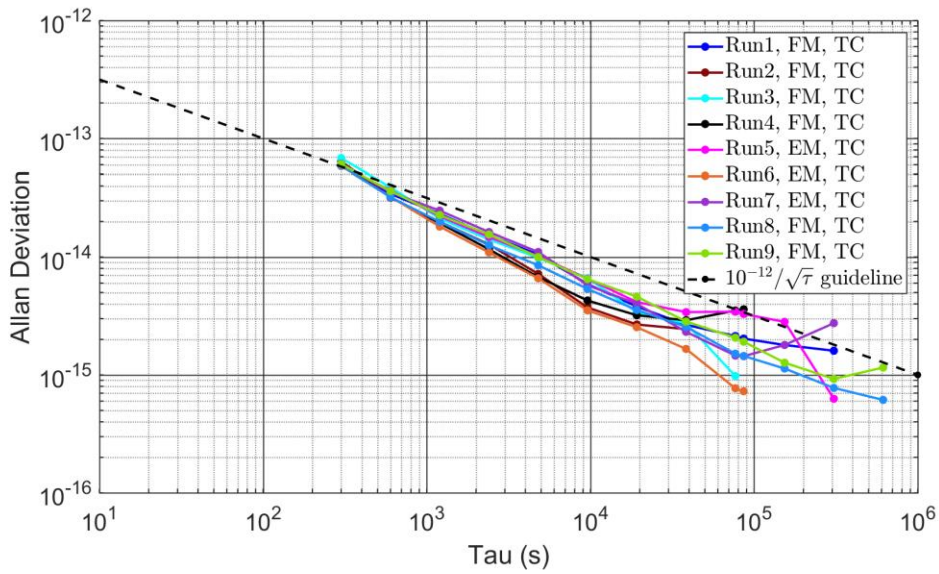

Figure 11. Single receiver (long-baseline) Allan deviations covering all the timespans investigated in this paper.

CONCLUSION

In conclusion, we have described features of a simple algorithm for treating data gaps and jumps in GPS carrier-phase data and for calculating the Allan deviation in the presence of these data gaps. We validated it and confirmed the DSAC EM and FM noise floor down to the $1\text{x}10^{-17}$ level (at $5\text{x}10^{5}$ s), by comparing two receivers sourced by the same H-maser clock at zero-baseline (Fig. 2). This noise floor is of course for an idealized environment -- temperature stable and zero-baseline -- but it shows that our technique for dealing with delay jumps and data gaps is not limiting us at that level, at least for these low noise common-clock data sets. In addition, we have shown evidence of being able to temperature compensate down to $3\text{x}10^{-17}$ (at $1\text{x}10^{5}$ s; Fig. 3). We have also overlaid multiple data sets to give a sense of the reproducibility of these noise floors, albeit, not an exhaustive study.

We plan to use this algorithm to monitor DSAC clock performance in space, which will be via single receiver data rather than the pair data we presented here. The Allan deviations will be higher due to long-baseline GPS time-transfer noise and/or clock noise. We plan to do further work to study whether the higher noise will mean jump corrections could introduce artifacts above the measured Allan deviation level. For now, our low zero-baselines achieved and our adequate long-baselines, $10^{-12}/\sqrt{\tau}$ at a day, give us confidence in using this algorithm and temperature correction for DSAC. Furthermore, on-orbit, single receiver noise floors have the potential to be better than on-ground single receiver performance, due to there being no on-orbit tropospheric fluctuations along different transmitter lines of sights. Finally, our zero-baseline pair noise-floors are a glimpse into the potential capability for future space-to-space comparisons (no atmosphere), but determining orbits will be a further challenge for achieving these low levels of noise in space.

APPENDIX

TABLE II
DATE RANGES OF DATASETS

Dataset	Date Range (GPS Time)	EM/FM Location
	Fig. 5	
Run1	12/21/2016 00:00 - 01/02/2017 23:59	FSTL
Run2	11/08/2015 00:00 - 11/08/2015 23:59	Bldg. 2, JPL
Run3	11/14/2015 00:00 - 11/15/2015 23:59	Bldg. 2, JPL
Run4	01/09/2015 00:00 - 01/11/2015 23:59	Bldg. 1, JPL
	Fig.6	
Run5	03/01/2017 05:50 - 03/08/2017 23:59	FSTL (EM in thermal chamber)
Run6	12/20/2014 02:00 - 12/22/2014 22:00	FSTL
Run7*	12/21/2016 00:00 - 12/28/2016 23:59	FSTL
Run8	11/23/2016 00:00 - 12/13/2016 23:59	FSTL
	Others	
Run9	02/23/2016 00:00 - 03/16/2016 23:59	Bldg. 1, JPL
Run10*	12/25/2016 0:00 - 12/28/2106 23:59	FSTL

* Different subsets of Run1

ACKNOWLEDGMENTS

We thank Angela Dorsey who developed the software for the DSAC GPS receiver, helped with the receiver testing, and helped with early versions of this manuscript. We also thank Rabi Wang who was in charge of equipment moves and installation in different buildings. Finally, we thank Bob Tjoelker for a critical review and suggestions on the manuscript and Shailen Desai who helped us implement the technique used to process zero/short-baseline data which removes common-mode troposphere and ionosphere delay.

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contracts with the National Aeronautics and Space Administration.

Copyright 2019 California Institute of Technology. U.S. Government sponsorship acknowledged.

REFERENCES

1. C. A. Greenhall, “A method for using a time interval counter to measure frequency stability,” *IEEE Trans. Ultrason., Ferroelectr., Freq. Control (TUFFC)*, vol. 36, no. 5, pp. 478-480, Sept. 1989.
2. C. A. Greenhall, A. Kirk, and R. L. Tjoelker, 2007, “A multi-channel stability analyzer for frequency standards in the Deep Space Network,” in *Proc. 38th Precise Time Time Intervall (PTTI)*, Reston, Virginia, USA, Dec. 2006, pp. 105-115.
3. G. Petit, A. Kanj, S. Loyer, J. Delporte, F. Mercier, and F. Perosanz “1×10^{-16} frequency transfer by GPS PPP with integer ambiguity resolution,” *Metrologia*, vol. 52, no. 2, pp. 301-309, Apr. 2015.
4. J. Yao, I. Skakun, Z. Jiang and J. Levine, “A detailed comparison of two continuous GPS carrier-phase time transfer techniques,” *Metrologia*, vol. 52, no. 5, pp. 666-676, Oct. 2015.
5. G. Petit and P. Defraigne, “The performance of GPS time and frequency transfer: comment on 'A detailed comparison of two continuous GPS carrier-phase time transfer techniques',” *Metrologia*, vol. 53, no. 3, pp. 1003-1008, Jun. 2016.

6. J. Leute, N. Huntemann, B. Lipphardt, C. Tamm, P. B. R. Nisbet-Jones, S. A. King, R. M. Godun, J. M. Jones, H. S. Margolis, P. B. Whibberley, A. Wallin, M. Merimaa, P. Gill, and E. Peik, "Frequency comparison of $^{171}Yb^{+}$ ion optical clocks at PTB and NPL via GPS PPP," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control (TUFFC)*, vol. 63, no. 7, pp. 981-985, Jul. 2016
7. S. Droste, C. Grebing, J. Leute, S. M. F. Raupach, A. Matveev, T. W. Hänsch, A. Bauch, R. Holzwarth and G. Grosche, "Characterization of a 450km baseline GPS carrier-phase link using an optical fiber link," *New J. Phys.*, vol. 17, no. 8, p. 083044, Aug. 2015.
8. A. Bauch, J. Achkar, S. Bize, D. Calonico, R. Dach, R. Hlavac, L. Lorini, T. Parker, G. Petit, D. Piester, K. Szymaniec and P. Uhrich "Comparison between frequency standards in Europe and the USA at the 10−15 uncertainty level," Metrologia, vol. 43, no. 1, pp. 109-120, Feb. 2006.
9. J. Ray and K. Senior, "IGS/BIPM pilot project: GPS carrier phase for time/frequency transfer and timescale formation," *Metrologia*, vol. 40, no. 3, pp. S270–S288, Jun. 2003.
10. G. Petit, C. Thomas, Z. Jiang, P. Uhrich, and F.Taris, "Use of GPS ASHTECH Z12T receivers for accurate time and frequency comparisons," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control (TUFFC)*, vol. 46, no. 4, pp. 941-949, Jul. 1999.
11. *DSN Telecommunications Link Handbook, Publication 810-005, Module 304 Rev. B*, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, Dec. 2014.
12. E. Burt, W. Diener, and R. Tjoelker "A compensated multi-pole linear ion trap mercury frequency standard for ultra-stable timekeeping," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control (TUFFC)*, vol. 55, no. 12, pp. 2586-2595, Dec. 2008.
13. R. L. Tjoelker, J. D. Prestage, E. A. Burt, P. Chen, Y. J. Chong, S. K. Chung, W. Diener, T. Ely, D. G. Enzer, H. Mojaradi, C. Okino, M. Pauken, D. Robison, B. L. Swenson, B. Tucker, and R. Wang "Mercury ion clock for a NASA technology demonstration mission," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control (TUFFC)*, vol. 63, no. 7, pp. 1034-1043, Jul. 2016.
14. T. A. Ely, T. Koch, D. Kuang, K. Lee, D. Murphy, J. Prestage, R. Tjoelker, J. Seubert, "The Deep Space Atomic Clock mission," in *Proc.23rd Int. Symp. Space Flight Dyn.*, Pasadena, CA, USA, Oct. 2012, pp. 1-15.
15. W. Bertiger, S. D. Desai, B. Haines, N. Harvey, A. W. Moore, S. Owen, and J. P. Weiss "Single receiver phase ambiguity resolution with GPS data," *J. Geod.*, vol. 84, pp. 327-337, Mar. 2010.
16. W. J. Riley, "Handbook of frequency stability analysis," NIST, U.S. Dept. Commerce, Boulder, CO, USA, Tech. Rep. 1065, Jul. 2008.
17. S. D. Desai, J. Gross, B. J. Haines, and D. A. Stowers, "Verification and Validation of the GNSS Stations at the Prototype Core Site for NASA's Next Generation Space Geodesy Network", Abstract G53B-0928 presented at 2013 Fall Meeting AGU, San Fransisco, CA, USA, Dec. 9-13.
18. J. Y. Tien, B. Bachman Okihiro, S. X. Esterhuizen, G. W. Franklin, T. K. Meehan, T. N. Munson, D. E. Robison, D. Turbiner, L. E. Young, "Next generation scalable spaceborne GNSS science receiver," in *Proc. of 2012 ITM of ION*, Newport Beach, CA, USA, Jan./Feb. 2012, pp. 882-914.