Timing and Location Performance of Recent u-blox GNSS Receiver Modules

by John Ackermann N8UR[1]

1. INTRODUCTION

In the past few years several companies have introduced small GNSS[2] modules intended for OEM timing and positioning applications. u-blox[3] AG, a Swiss corporation, is perhaps the most well-known of these and has introduced several generations of receivers with increasing capabilities. This paper presents experimental data on the seven u-blox modules listed in Table 1.[4] More detailed information about the capabilities of the units is contained in Appendix 1.

Model	Comments
LEA-M8F	Frequency and timing series (disciplined frequency source)
NEO-M8N	Navigation series (low cost)
NEO-M8P	Positioning series (internal RTK engine)
NEO-M8T	Timing series (dual timepulse)
NEO-M9N	Navigation series (low cost, still L1 only)
ZED-F9P	Positioning series (L1/L2, internal RTK engine)
ZED-F9T	Timing series (L1/L2, dual timepulse)

Table 1: u-blox modules tested.

The focus of this paper is on the receivers' timing performance, and primarily the stability and other characteristics of its hardware time pulse output. Section 2 describes the timekeeping performance of the receivers. Sections 3 through 6 explore other aspects of timekeeping performance. Section 7 briefly explores how the performance of these receivers can allow a different design philosophy for GPS disciplined oscillators ("GPSDOs"). Finally, Section 8 provides a limited overview of the receivers' positioning performance.

1 The author has no relationship with u-blox AG.

2 These are "Global Navigation Satellite System" or "GNSS" receivers because they support satellite constellations in addition to the United States' Global Positioning System. However, in this paper the term "GPS" is used because, as of this writing, in most timing applications only the GPS constellation is used and the U.S.N.O. master clock serves as the reference. For precise positioning applications, both GPS and to a lesser extent the Russian GLONASS systems are used. The Chinese BeiDou and European Galileo systems are just starting to reach operational use.

3 Based on the corporate web site, the correct name is "u-blox" with hyphen and without capitalization.

4 The author acknowledges and greatly appreciates support from NSF Grants AGS- 2002278 and AGS-1916690, The University of Scranton, and the New Jersey Institute of Technology Center for Solar-Terrestrial Research.

1.2 Testing methodology

For timing purposes, the key criterion is the quality of the receiver's hardware output that provides an electrical pulse aligned to GPS time. This is usually, but not always 1 pulse per second ("PPS"). U-blox calls this the TIMEPULSE output. The seven receivers were simultaneously and individually evaluated by using a counter to measure the offset in time between their TIMEPULSE outputs and a local PPS signal derived from a stable atomic clock. The variations in the second-by-second values recorded show the noise (sometimes referred to as "jitter") of the timing signal, expressed as quantities of time.[5]

All seven GPS modules, as well as a CNS Systems CNS-Clock II GPS time receiver[6] used for sanity-checking, were fed from the same dual-frequency GPS antenna (Trimble Zephyr Geodetic[7]) via distribution amplifiers. Their TIMEPULSE outputs were measured simultaneously using a multi-channel counter (TAPR multi-TICC[8]) which has eight independent input channels with resolution of about 60 picoseconds. A high-performance Cesium frequency standard (HP 5071A[9]) served as the reference clock driving the multi-TICC counter. The measurement campaign lasted just under six days and recorded more than 500,000 samples from each of the eight receivers.

The series of timestamps recorded from each receiver's TIMEPULSE output constitutes a record of the time offset, or phase, of that output relative to the reference clock. The timestamp data was captured to text files and processed with the widely-used TimeLab[10] time and frequency analysis software written by John Miles. The figures presented below, unless otherwise stated, were rendered by TimeLab after the applicable phase records were loaded and processed.

For some of the measurements reported in this paper, additional data collection runs were used to enable testing of various configurations.

Unless otherwise stated, for all tests the receivers were set to their default configurations save only for choosing the 0-D (timing) solution mode where available. No compensation was made for cable or other systematic delays, and no attempt was made to measure absolute time accuracy. Receivers that allowed entry of fixed observation coordinates were set to ECEF[11] X, Y, and Z values previously derived from post-processed Precise Point Positioning measurements made with a Trimble NetRS[12] receiver using the same antenna as the one used for these measurements.

---

5 For example, PPS noise might be described as "20 nanoseconds RMS". Somewhat confusingly, it is common to refer to the *amplitude* of noise; it is important to realize that this refers to the magnitude of time variations and not voltage or strength.

6 https://www.cnssys.com/

7 https://kb.unavco.org/kb/article/the-design-and-performance-of-the-zephyr-geodetic-antenna-trimble-publication-241.html. Note that this antenna is not optimized for GLONASS or BeiDou frequencies; however it appears to receive GLONASS signals reasonably well, and BeiDou performance was not tested.

8 https://github.com/TAPR/TICC/blob/master/multi-ticc/multi-TICC_App_Note_2020-01.pdf

9 Currently sold by Micro-Semi division of MicroChip: https://www.microsemi.com/product-directory/cesium-frequency-references/4115-5071a-cesium-primary-frequency-standard

10 http://www.ke5fx.com/timelab/readme.htm

11 "Earth-Centered, Earth Fixed" coordinates in meters

12 https://kb.unavco.org/kb/article/trimble-netrs-resource-page-471.html

1.3 Interpreting the Results

The time pulse output of a GPS receiver typically runs at a one pulse per second rate, and over the long term tracks the master clock of the GPS satellite constellation, which in turn is traceable to the US. Naval Observatory and National Institute of Standards and Technology time and frequency standards. In other words, a GPS receiver with a timing output provides a replica of the official time, and because frequency can be derived from a series of time measurements, of standard frequency as well. However, in the short term these pulses contain noise, or "jitter," that results from both limitations in the quality of the GPS solution obtainable, and limitations in the hardware capability of the module. The characteristics of that noise determine the short-term timing capabilities of the unit.

It is possible to plot the PPS values on a graph with the Y axis showing the value and the X axis showing the elapsed time – what is called a "strip chart" recording. This technique gives a qualitative view of the data, and from it one can obtain a sense of performance. However, such a presentation provides mainly a qualitative view and does not lend itself to any but the most basic quantitative evaluation. It is often helpful to look at both the full phase record, as well as a close-up view that shows short-term (second-by-second) changes. This can reveal performance characteristics that are hidden by the limited resolution of a plot showing thousands of data points.

The noise of a clock or oscillator[13] can be analyzed statistically to help understand the stability of the clock's tick over varying periods of time. For example, the standard deviation of the series of PPS values could be used to get an idea of their spread.

However, for reasons beyond the scope of this paper, standard deviation is not the best tool to analyze the noise processes at work in clocks. A related statistic called the Allan Deviation ("ADEV") is designed specifically for frequency stability analysis and better serves the purpose.[14] In very general terms, ADEV can be considered as the likely variation between any two measurements taken at a specified interval (the interval is denoted as "tau"). For example, stating that "the "Allan Deviation of the PPS output is 2.3×10^{-10} at tau = 10 seconds" means that the values of any measurements taken of this PPS source at intervals of 10 seconds will mainly be within a range of 23 nanoseconds. A table of ADEV vs. tau describes the stability of the oscillator over varying time periods.

It is convenient to plot ADEV versus tau on a graph with ADEV on the Y axis, and tau on the X axis, both in log format. One advantage of this representation is that the slope of the plot reveals the primary noise process at work in that range of tau, as shown in Figure 1.

13 While "clock" and "oscillator" have different formal definitions (a clock consists of an oscillator plus additional components), this paper follows common informal practice and uses the two terms synonymously.

14 A quite good tutorial on ADEV is at https://en.wikipedia.org/wiki/Allan_variance

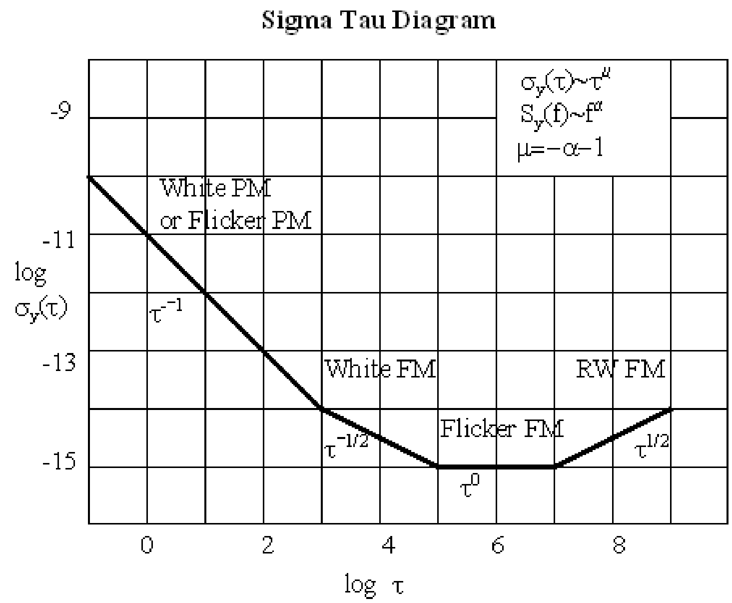

Figure 1: Noise Processes Shown as Allan Deviation Slope[15]

The PPS output of most GPS receivers exhibits white phase modulation bounded in absolute amplitude except for outliers, so the ADEV normally improves by one order of magnitude for each order of magnitude increase in tau. This is shown by a slope of minus 1 on the ADEV plot. In other words, longer averaging continues indefinitely to reduce noise and increase frequency stability (and thereby allow a more precise measurement).

In the discussion below, both strip-chart phase records and ADEV plots are used for illustration. Often both the phase record of the full data run, as well as a zoomed-in portion showing second-by-second variations, are provided.

As a technical note, the phase record will show any offset in frequency between the reference clock and the device under test as an upward or downward trend whose slope can be converted to a fractional frequency offset value. Since no two clocks ever run at exactly the same frequency, if for no reason other than quantum uncertainties, such an offset will always appear in the phase record of two independent sources. The offset and slope between the Cesium frequency reference and GPS constellation clock can be seen in Figure 3 below. For the purposes of this paper, the frequency offset is not relevant to the receiver performance, and therefore all figures showing phase plots, other than Figures 3 and 34, have had this offset removed to present a flat phase trend. Note that Allan Deviation measurements are not sensitive to frequency offset, so there is no need to account for it in ADEV plots.

15 Figure courtesy of W. J. Riley

2. RESULTS: TIMING PERFORMANCE

To set the stage, the following figures show the timing performance of all seven receivers in a single set of plots.

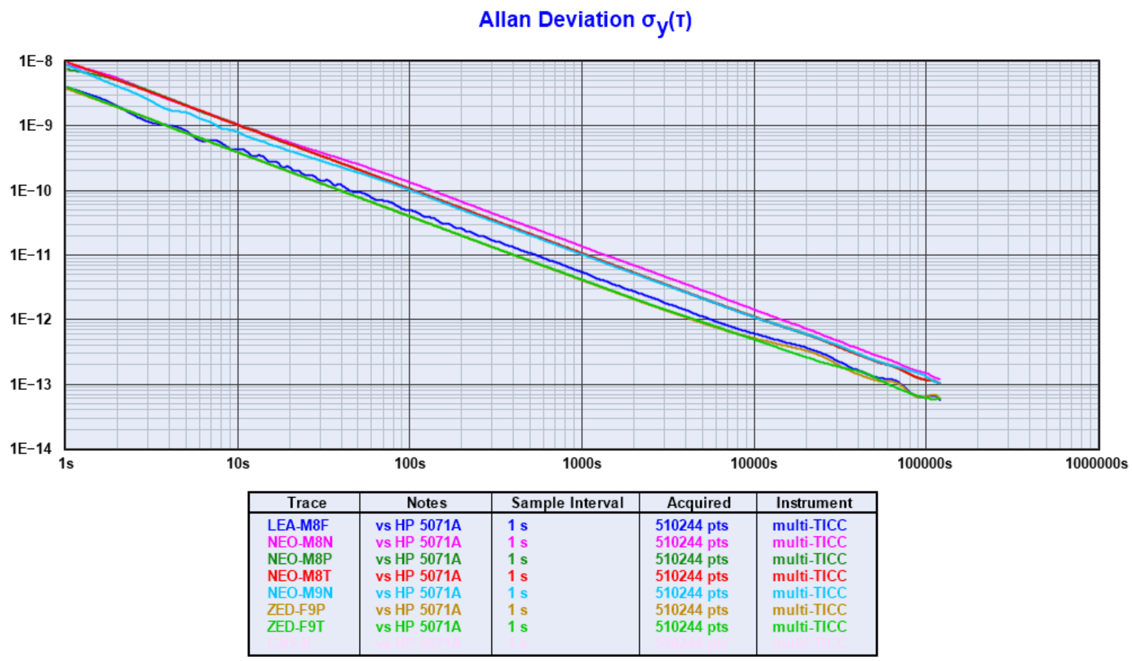

Figure 2: Allan Deviation of all seven u-blox receivers.

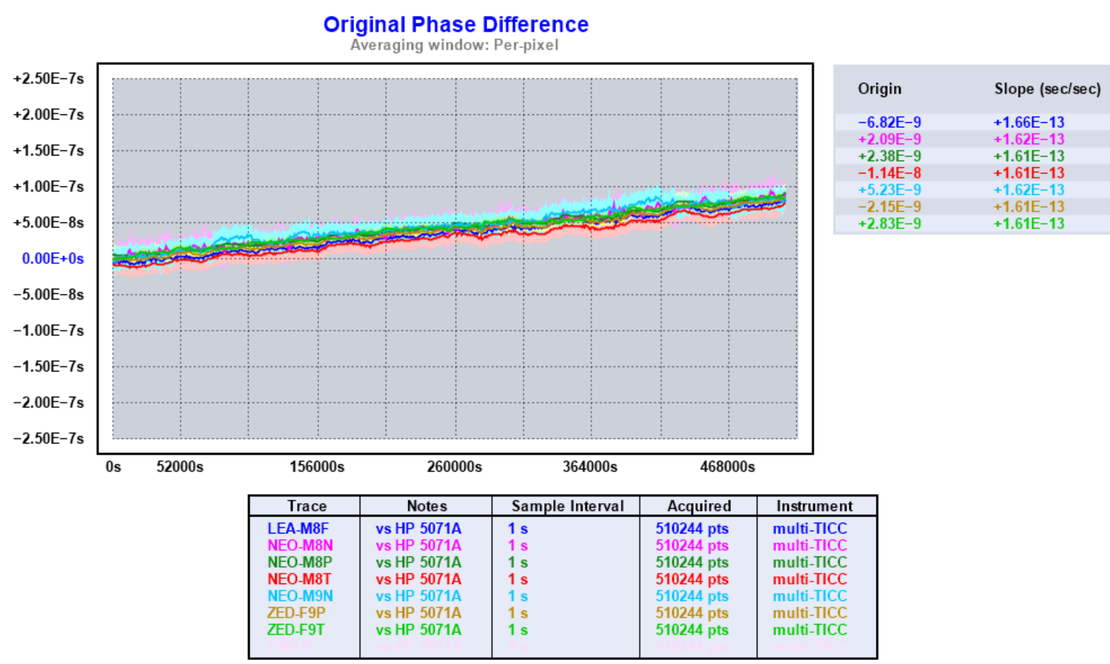

Figure 3: Raw Phase Difference

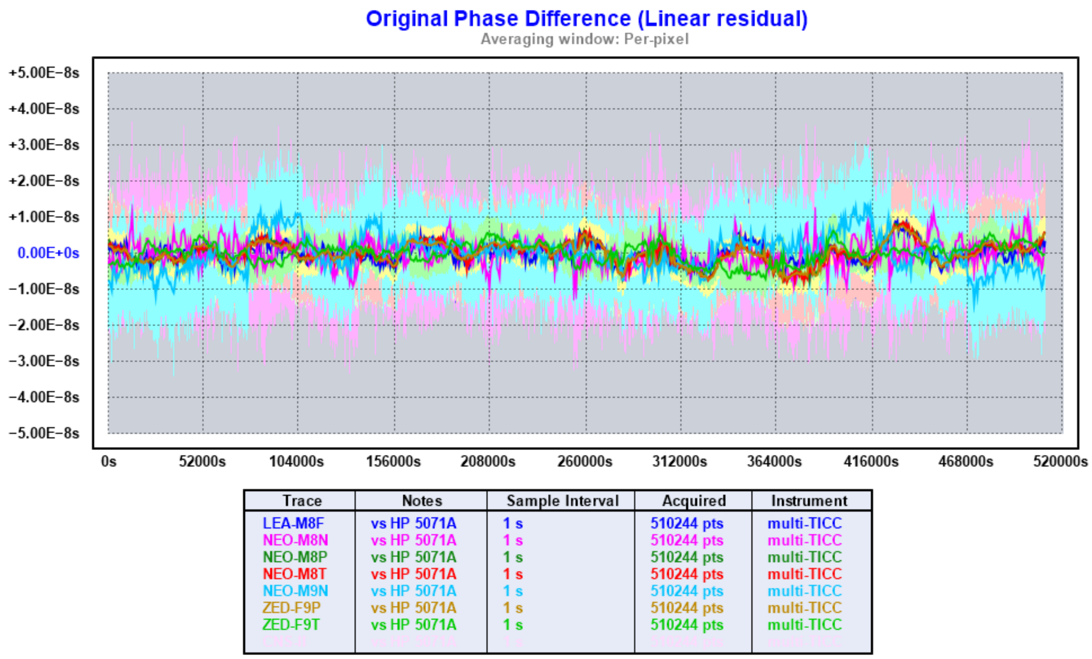

Figure 4: Phase With Offset Removed

Figure 2 shows the Allan Deviation of the seven u-blox receivers on a single plot. The results fall into two main groups: the single-frequency receivers with ADEV around 1×10^{-8} at 1 second tau, and the dual-frequency receivers, as well as the LEA-M8F, with 1 second ADEVs near 4×10^{-9}.

Figure 3 shows the phase of the receivers compared to the Cesium reference. The upward slope indicates that the reference was about 1.6×10^{-13} low in frequency compared to the GPS constellation.[16] Figure 4 shows the same data with that slope removed.

Because it is difficult so see the performance of the individual receivers in these composite plots, the following sections describe and plot results for subgroups of the receivers.

16 The HP 5071A/HP specification is for frequency accuracy of better than 5×10^{-13}; this measurement shows that this unit is operating well within that specification.

2.1 “N” Series Receivers

The NEO-M8N and NEO-M9N are low cost modules intended for navigation. They are not optimal for timing purposes because they do not allow a “0D” timing solution configuration (see Section 4 below), and they do not report the quantization error (see Section 3 below) to allow for software correction.

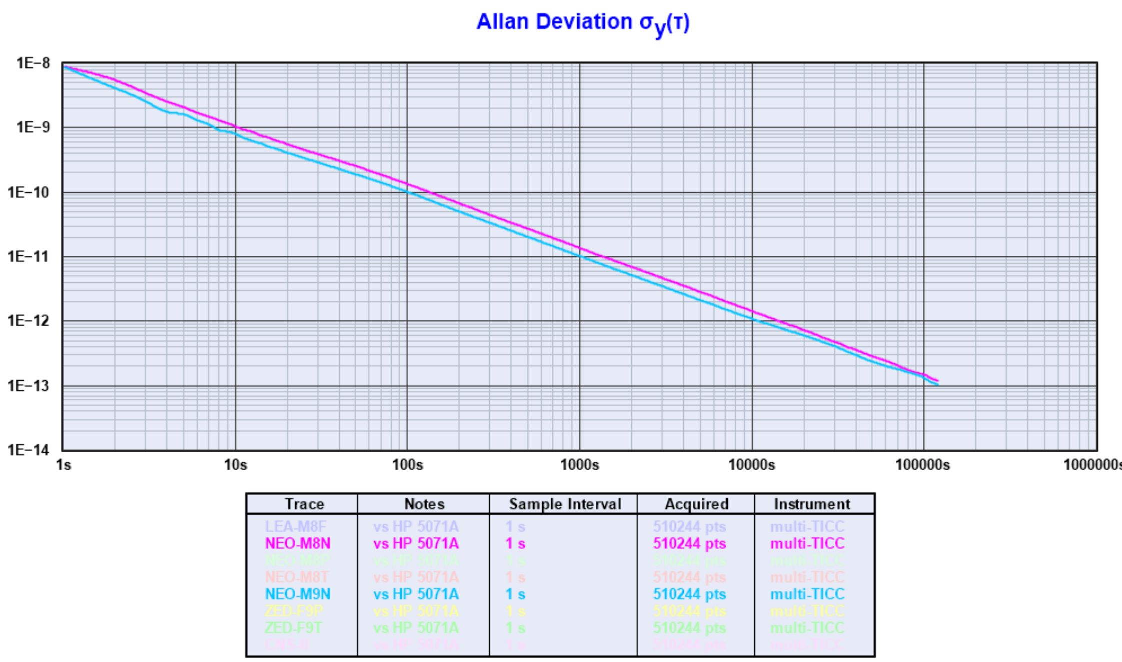

Figure 5: Allan Deviation of M8N and M9N Receivers

Figure 5 shows that the newer M9N receiver offers slightly better ADEV at most tau than the older M8N, but the difference is not substantial.

Figure 6 shows phase plots of the two receivers over the full measurement, and Figure 7 shows an approximately 45 second interval of that data. Figure 7 shows that the M8N has a sawtooth characteristic with a period of just under 10 seconds,[17] with a peak-to-peak range of about 17 nanoseconds. The M9N shows a different pattern, with noticeable second-to-second variation of about 7 nanoseconds and larger steps of about 14 nanoseconds at an approximate 9 second interval. The smaller average noise amplitude explains the slight ADEV advantage of the M9N unit.

17 There is no reason to believe that this period is consistent across receivers or measurement runs. In particular, the period may be temperature dependent.

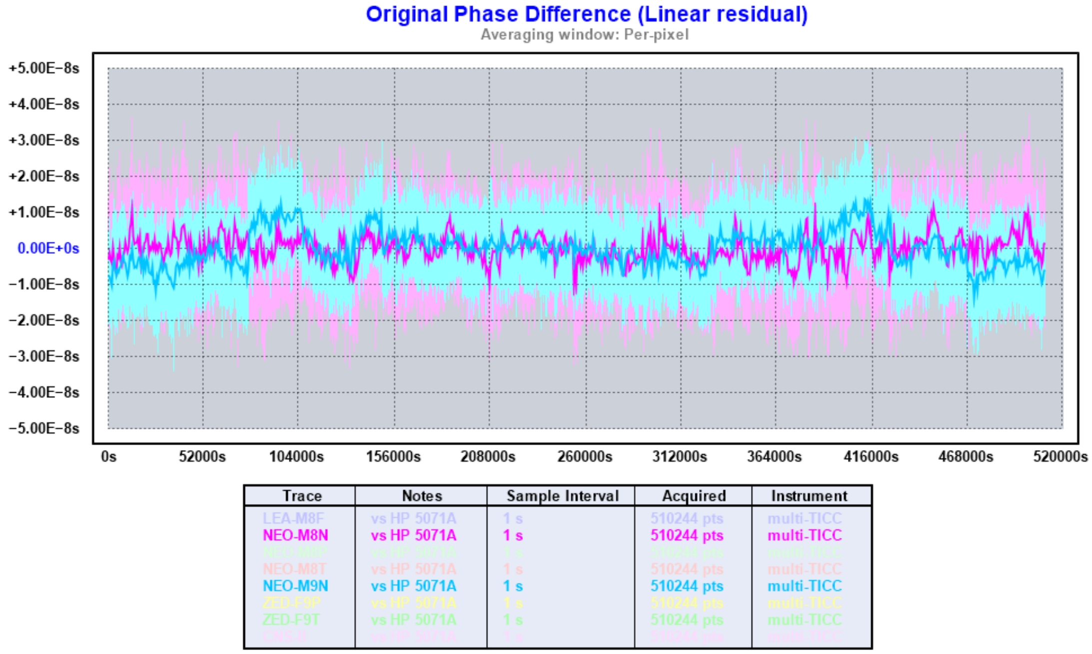

Figure 6: Phase Plot of M8N and M9N Receivers

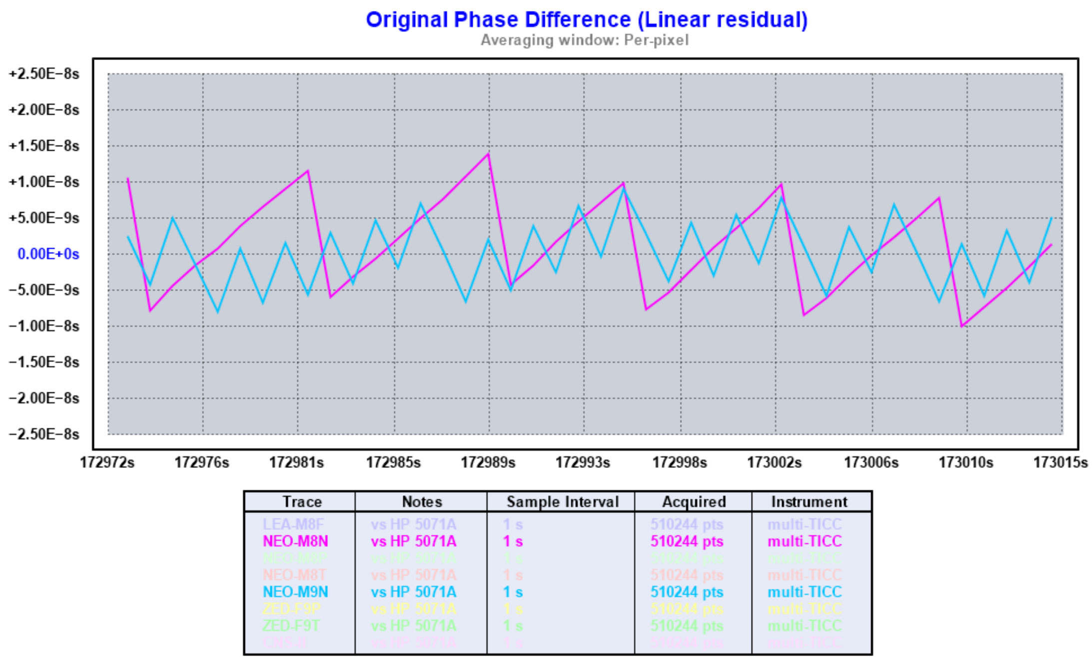

Figure 7: Zoomed Phase of M8N and M9N Receivers

2.2 M8P and M8T Series

The NEO-M8P (Positioning) and NEO-M8T (Timing) are more capable receivers than the “N” series, albeit still single-frequency. They are very similar, and while their raw timekeeping performance is only moderately better than the “N” series receivers, their additional capabilities allow better ultimate results. Both receivers can output the raw observation data (pseudorange, carrier phase, and doppler) required for RTK or PPP processing, and both provide quantization error correction (see Section 3 below), and a “0D” or “timing” solution mode that improves timing performance (see Section 4 below).

The differences between the “P” and “T” versions lie in the additional capabilities each provides. In particular, the M8P has an inbuilt RTK engine for real time kinematic positioning,[18] while the M8T does not include the RTK engine but provides two TIMEPULSE outputs and two EXTINT inputs versus one of each on the M8P. The M8T is somewhat less expensive than the M8P, and that it makes it an attractive choice where internal RTK processing is not required.

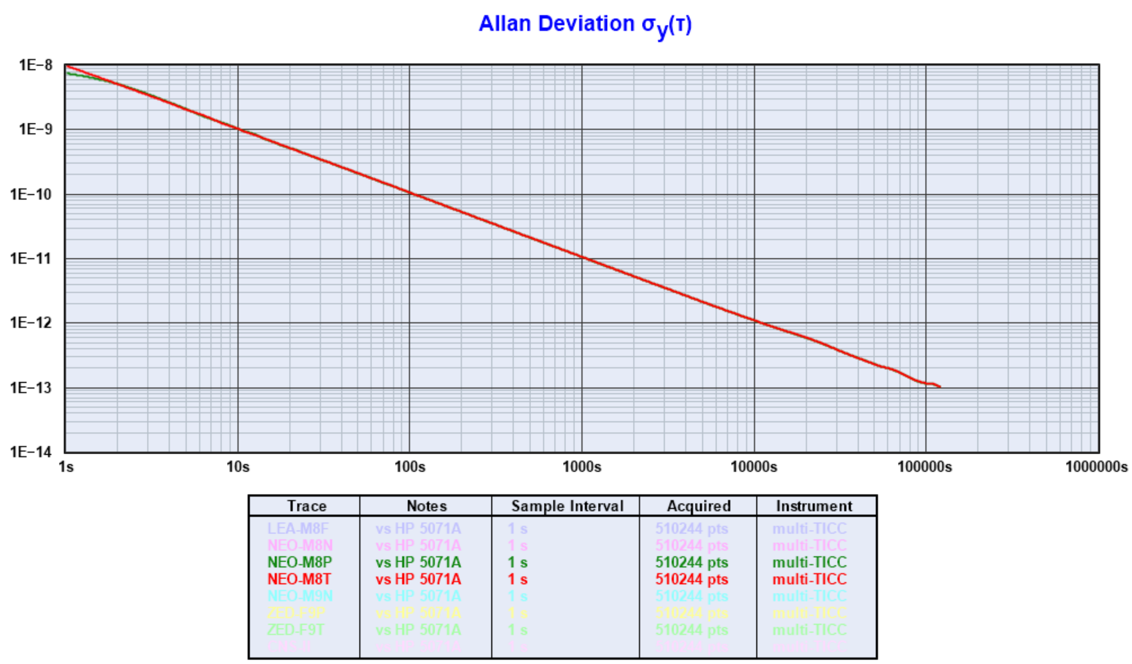

Figure 8: Allan Deviation of M8P and M8T Receivers

Figures 8 and 9 show ADEV and phase outputs for both the M8P (green) and M8T (red). These plots of the two receivers are virtually identical, with a very tiny edge to the M8P right at 1 second tau. The large-scale phase variations track very closely between the two units. Figure 10 zooms in the phase view. The peak-to-peak amplitude of both receivers is very similar, but the M8T sawtooth rate is about double that of the M8P and seems to operate in a “two-step” fashion (easier to see in the figure than to describe). The reason for this is currently unknown.

18 The RTK capability uses a second communications port to receive corrections in RTCM format, which are then applied to the navigation solutions output on the primary port.

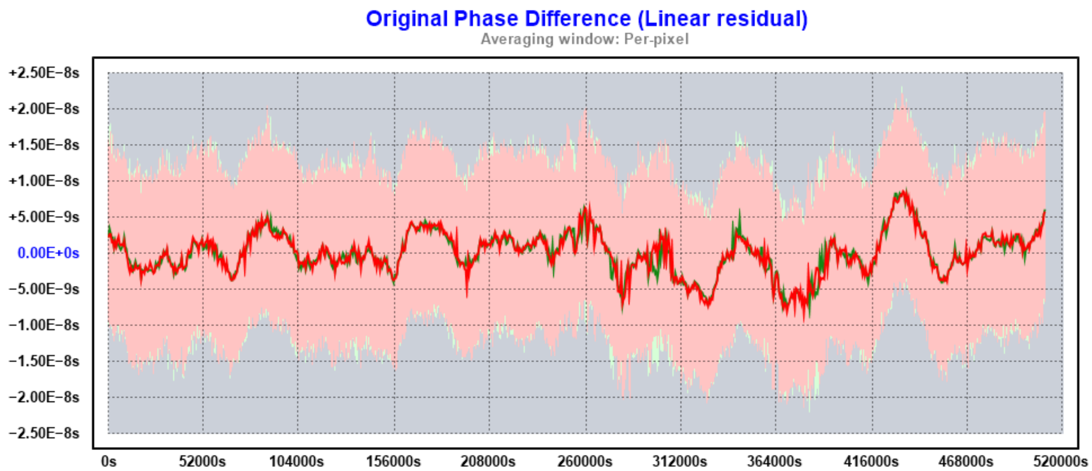

Trace	Notes	Sample Interval	Acquired	Instrument
LEA-M8F	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8N	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8P	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8T	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M9N	vs HP 5071A	1 s	510244 pts	multi-TICC
ZED-F9P	vs HP 5071A	1 s	510244 pts	multi-TICC
ZED-F9T	vs HP 5071A	1 s	510244 pts	multi-TICC

Figure 9: Raw Phase of M8P and M8T Receivers

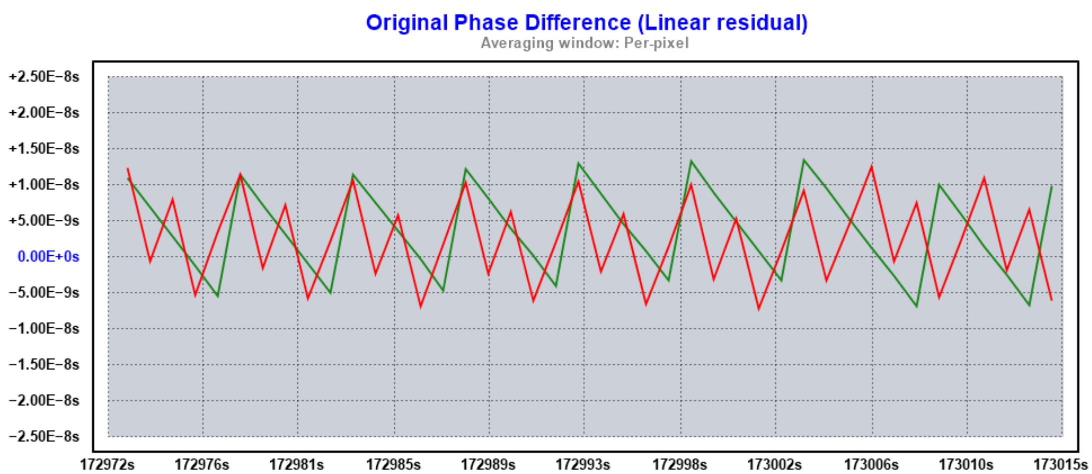

Trace	Notes	Sample Interval	Acquired	Instrument
LEA-M8F	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8N	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8P	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M8T	vs HP 5071A	1 s	510244 pts	multi-TICC
NEO-M9N	vs HP 5071A	1 s	510244 pts	multi-TICC
ZED-F9P	vs HP 5071A	1 s	510244 pts	multi-TICC
ZED-F9T	vs HP 5071A	1 s	510244 pts	multi-TICC

Figure 10: Zoomed Phase of M8P and M8T Receivers

2.3 F9P and F9T Series

The recently released ZED-F9P and ZED-F9T receivers represent a first in low-cost GPS technology, providing dual frequency operation along with other performance improvements. As in the “8” series, the ZED-F9P is optimized for positioning and has an inbuilt RTK processing engine, while the ZED-F9T is optimized for timing and costs less. Both provide raw observation data information for external processing.

Use of two reception frequencies (L1 = 1575 MHz; L2 = 1242 MHz) allows the receiver to calculate ionospheric/atmospheric delays and compensate for them, improving the accuracy of the fix. Survey/geodetic grade GPS timing receivers all utilize this strategy, and this capability should allow a significant performance increase over single-frequency receivers. It is also possible from two-frequency observations to calculate Total Electron Content values for the atmosphere near the receiver, a quantity of interest to space scientists and propagation predictors.

The dual-frequency capability of the ZED-F9 receivers has one qualification that is important to note: the only L2 modulation supported for the GPS constellation is the L2C signal that has been included in new GPS satellites launched since 2005.[19] For older satellites, the F9 receivers will revert to single-frequency operation. A few of the older satellites are still operational and they act to “dilute” the F9 performance compared to other dual-frequency receivers that can make use of additional L2 signals.

As a result, an F9 series receiver listening only to the GPS constellation may have marginally greater ambiguity in its results compared to a survey-grade receiver. As time goes on and new satellites replace the oldest ones, the importance of this limitation will diminish. In the short term, enabling the GLONASS constellation will increase the number of L2 satellites the F9 can use, and improve its positioning performance. However, for timing purposes it is recommended to use only one constellation,[20] so in that case using GPS and disabling GLONASS is still the best option.

19 Early dual-frequency receivers used the L2P(Y) code which is broadcast by all GPS satellites. Though the code is encrypted and theoretically available only to military and other authorized users, means have been found to use the P code without decryption.

20 Because the two constellations are referenced to different master clocks, and mixing them increases timing ambiguity.

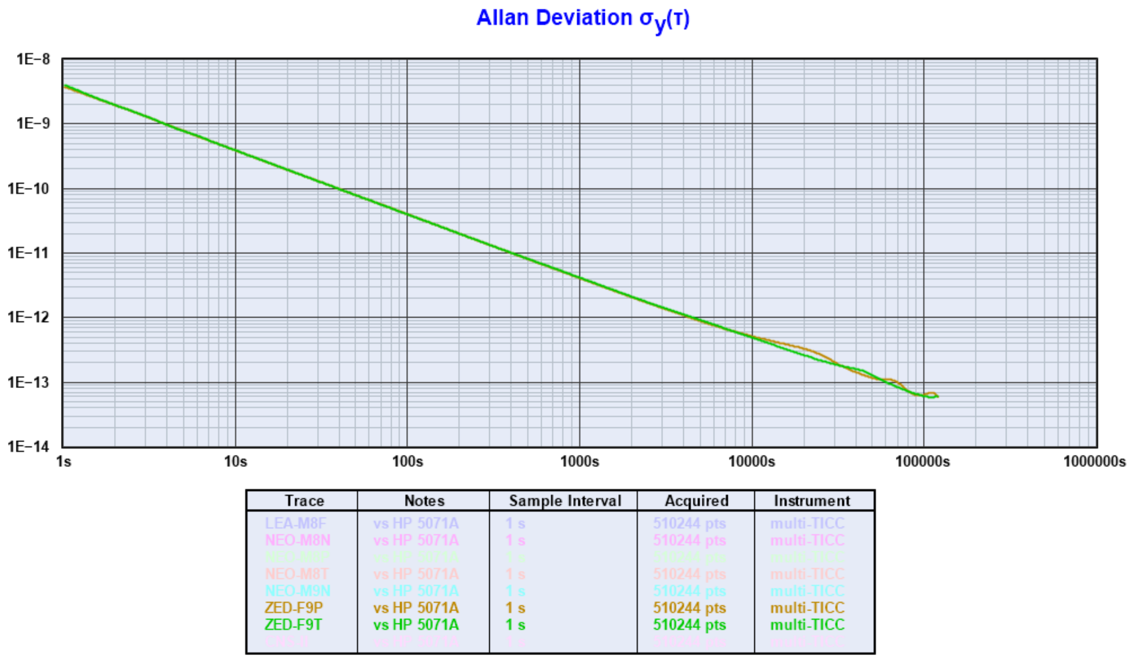

Figure 11: Allan Deviation of ZED-F9P and ZED-F9T Receivers

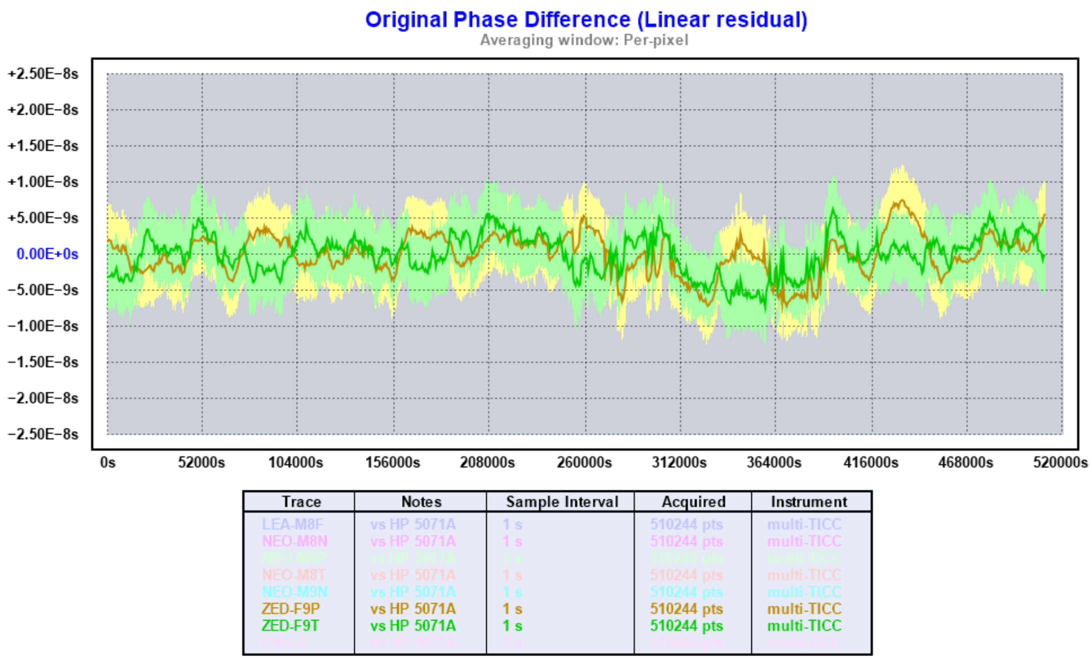

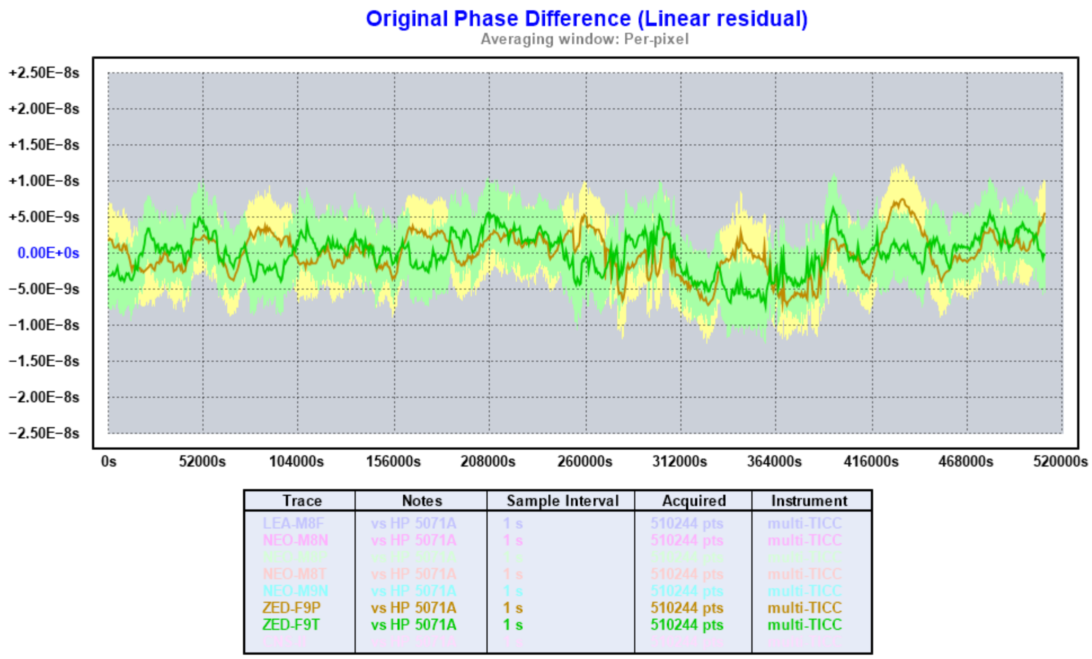

Figure 12: Raw Phase of ZED-F9P and ZED-F9T Receivers

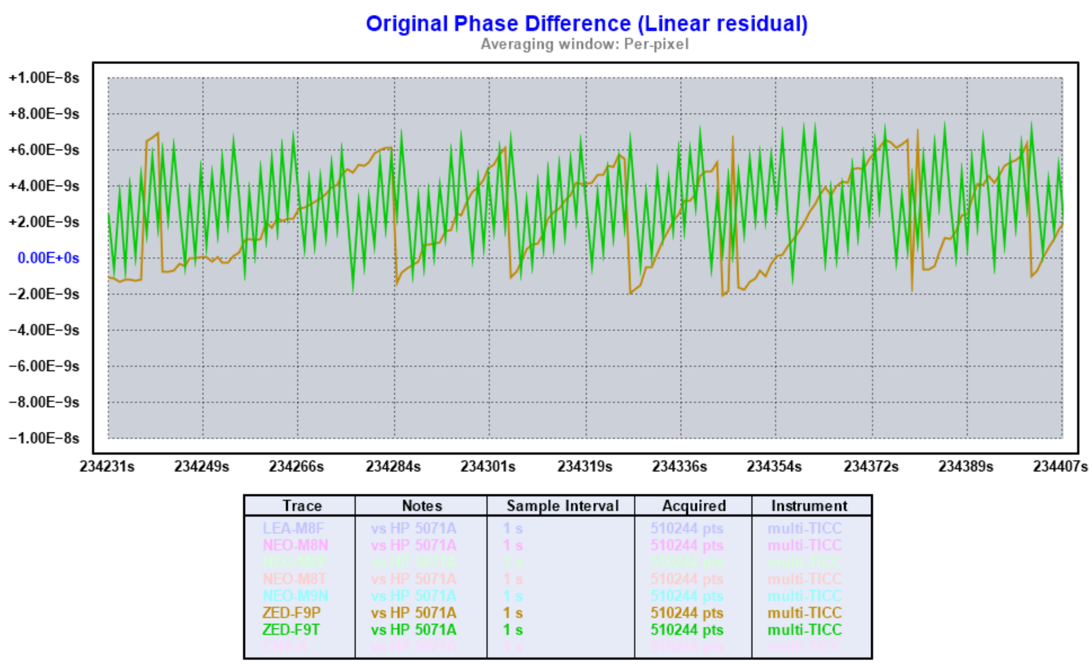

Figure 13: Zoomed Phase of ZED-F9P and ZED-F9T Receivers

Figures 11 and 12 plot ADEV and phase for the F9P and the F9T. As with the M8P and M8T, the two variants show essentially identical performance. It is interesting to note, but probably not meaningful, that in the long term the phase offsets of the two receivers seem to be 180 degrees out of phase – when one exhibits a positive offset, the other tends to show a negative offset.

Figure 13 shows an expanded phase view, and demonstrates markedly different behavior compared to the M8 series. In particular, the two receivers have very similar peak-to-peak noise of about 7 nanoseconds, but their processing appears to handle the TIMEPULSE output steering very differently. The F9T shows a multiple sawtooth pattern with second-by-second jitter of about 3 nanoseconds superimposed on an approximately 8 nanosecond sawtooth with a period of several seconds, while the F9P appears to show only a single sawtooth that occurs less frequently but with larger steps. As in the case of the 8-series receivers, the reason for this apparent difference is unknown.

2.4 Timing/Positioning Receiver Comparison By Series

Given the similar performance within each usage category (timing and positioning), it is useful to view the 8 and 9 series timing and positioning receivers as two groups, eliminating the lower-performing “N” series receivers and the interesting but quite different LEA-M8F from the visible plots.

For convenience, Figures 14 through 16 compare the four higher-end receivers (NEO-M8P, NEO-M8T, ZED-F9P, ZED-F9T) on a common scale. As expected, the dual-frequency receivers show lower jitter than their single-frequency equivalents, translating to lower Allan Deviation at a given measurement interval.

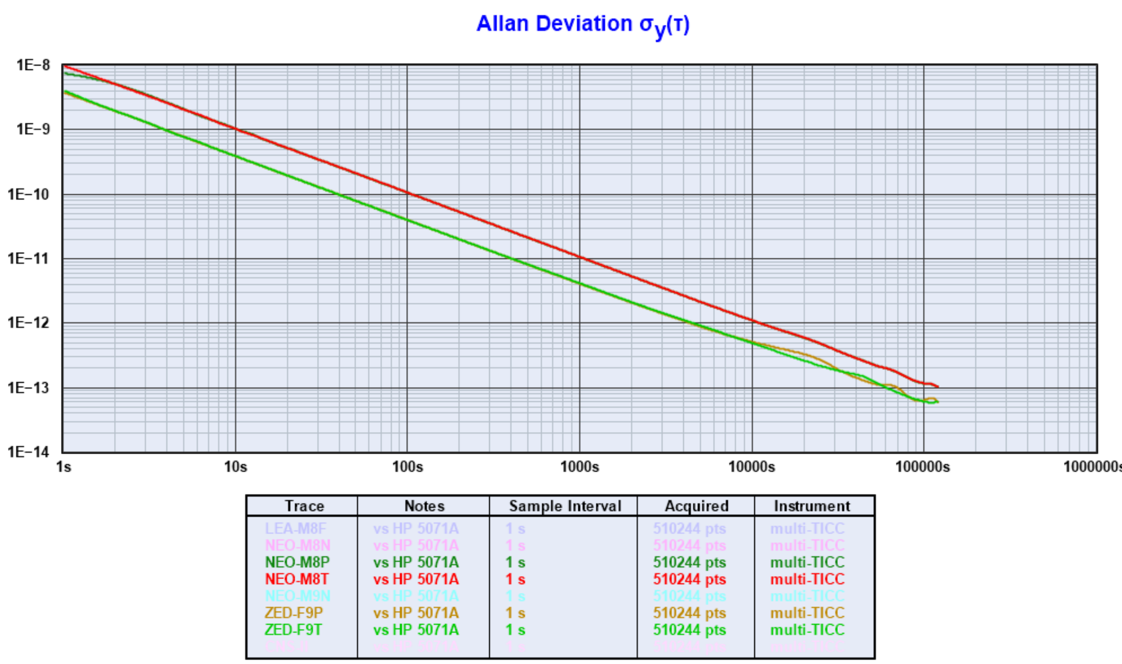

Figure 14: Positioning/Timing Receiver Allan Deviation Comparison

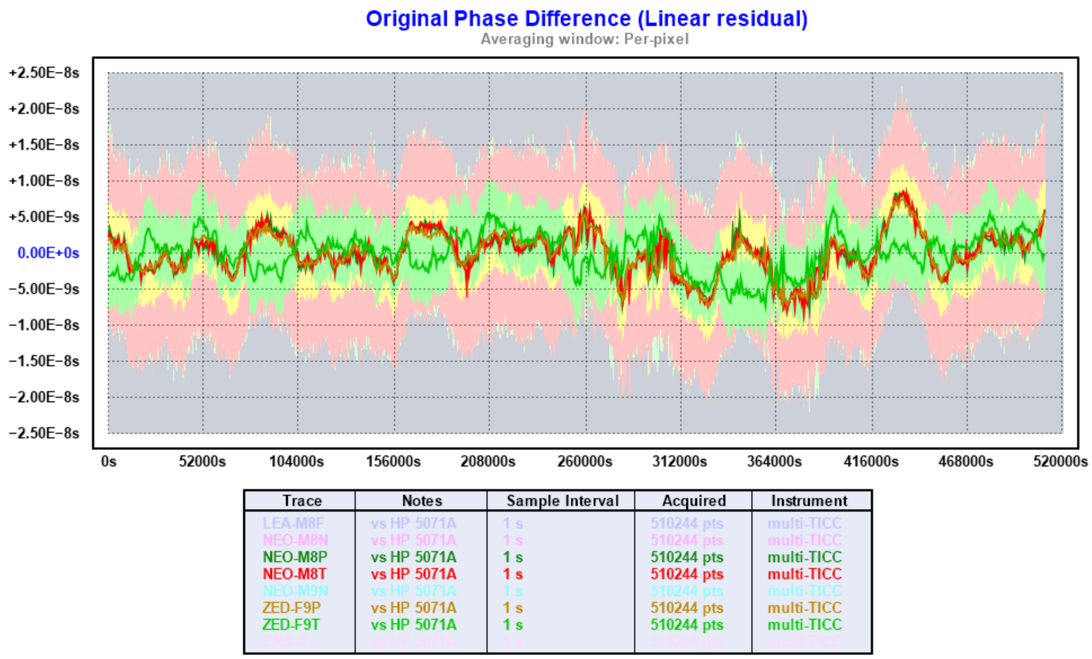

Figure 15: Positioning/Timing Receiver Phase Comparison

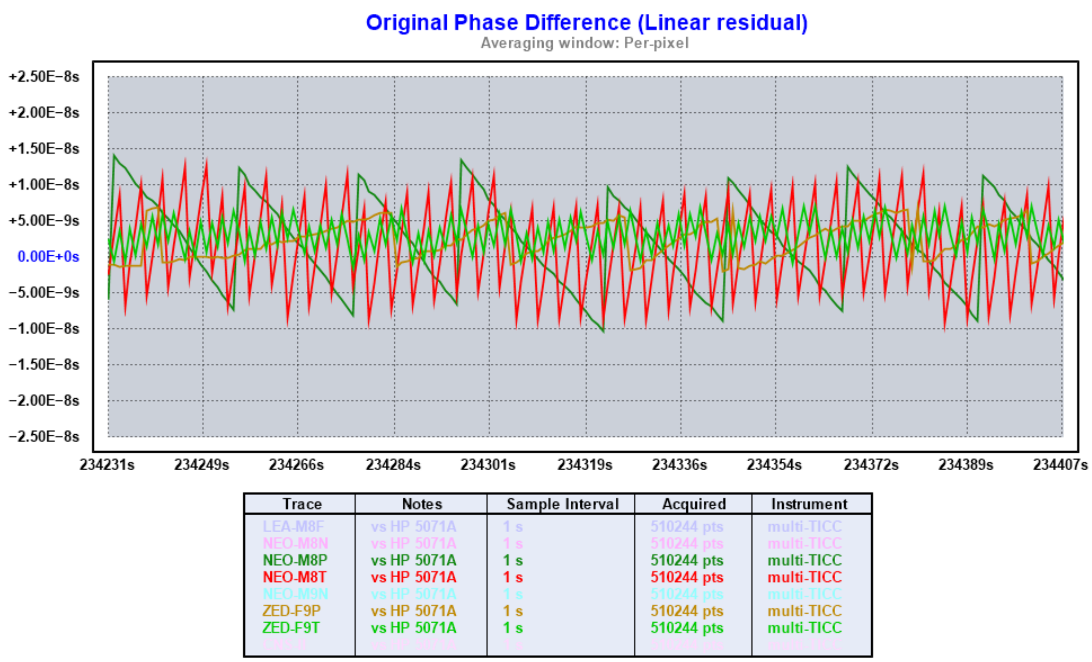

Figure 16: Positioning/Timing Receiver Zoomed Phase Comparison

The zoomed phase shows the two “P” receivers exhibiting the same simple sawtooth behavior over multiple seconds, while the two “T” models show a shorter-term sawtooth that is not visible on the “P” unit plots. It is unknown, but worth exploring in future, whether those differences are coincidental or result in differences in implementation.

2.5 LEA-M8F

The LEA-M8F has a different architecture than the other u-blox receivers. It includes a "frequency and time" subsystem that incorporates a 30.72 MHz TCXO which is kept disciplined to GPS. The TIMEPULSE is derived from that signal, and u-blox claims that it is essentially "jitter free". (Of course, the truth of this statement depends on the user's definition of jitter!)

In a sense, the LEA-M8F is a GPS disciplined oscillator ("GPSDO"), with its internal TCXO steered to the GPS timebase. In addition, it can be configured to steer an external oscillator. However, critical control loop parameters, such as time constant, are not made available to the user, so there is little opportunity to optimize its GPSDO performance, and in particular to take advantage of an oscillator with better short-term performance than the internal unit.

For this receiver, the hardware TIMEPULSE exhibits much less noise both in terms of short term jitter and GPS noise than other receivers in the M8 series – in fact, its performance is quite similar to the ZED-F9. However, the LEA-M8F does ***not*** include a usable quantization error message as discussed in Section 3 below), so the TIMEPULSE quality cannot be improved further by software correction.

The LEA-M8F makes its 30.72 MHz TCXO output available. This report includes measurements of that signal as well as the TIMEPULSE output.

2.5.1 LEA-M8F Timing Performance

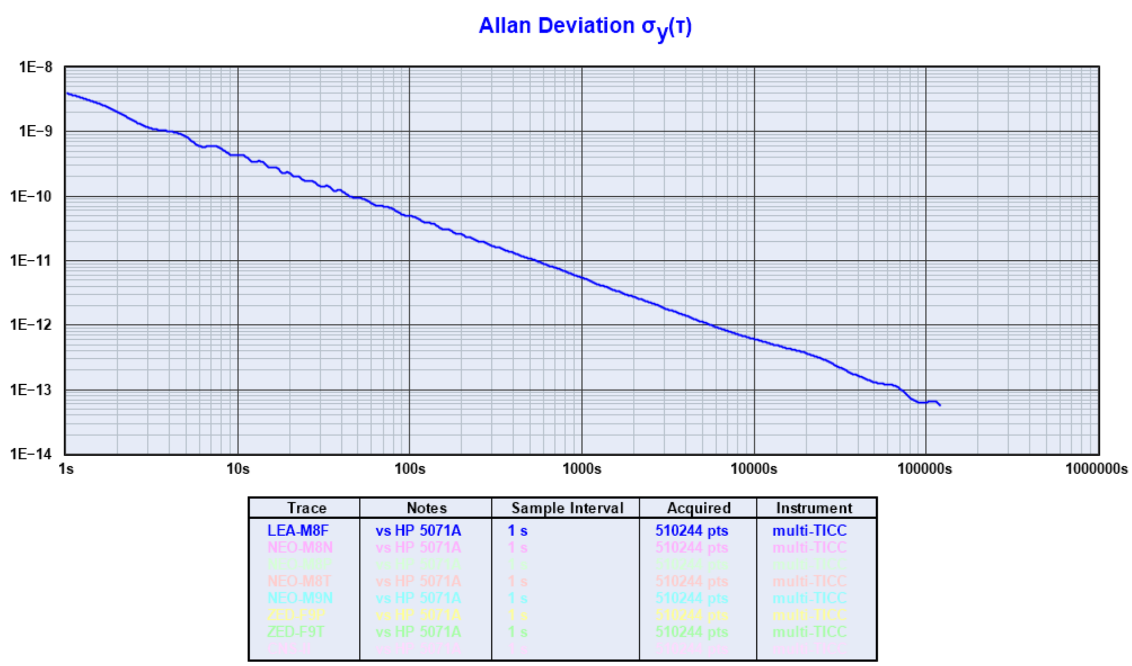

Figure 17: Allan Deviation of LEA-M8F Receiver

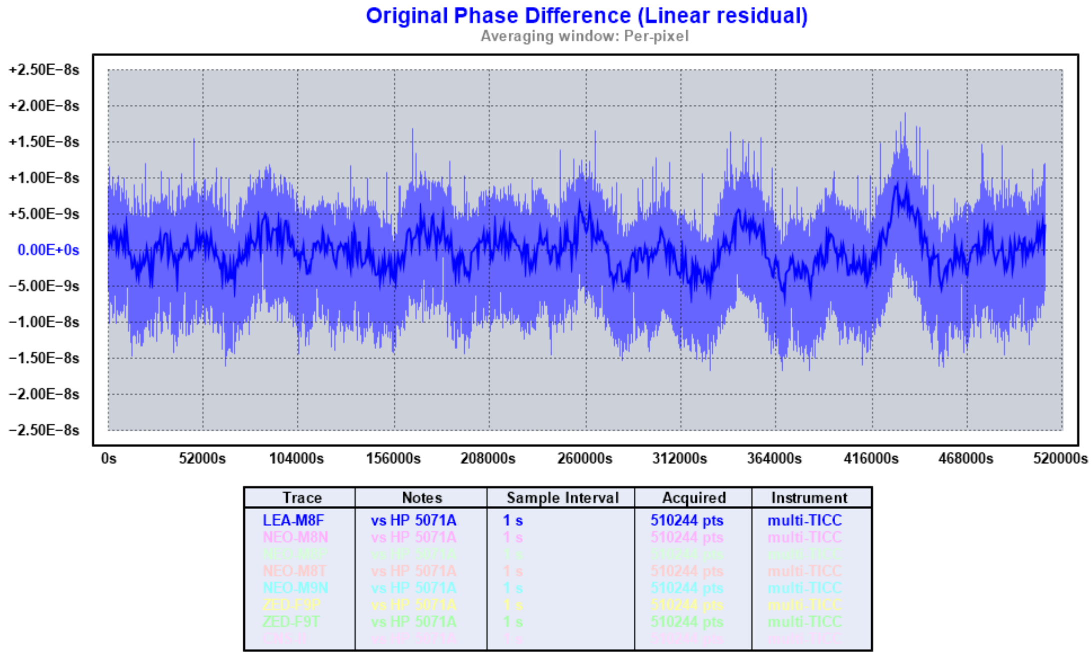

Figure 18: Raw Phase of LEA-M8F Receiver

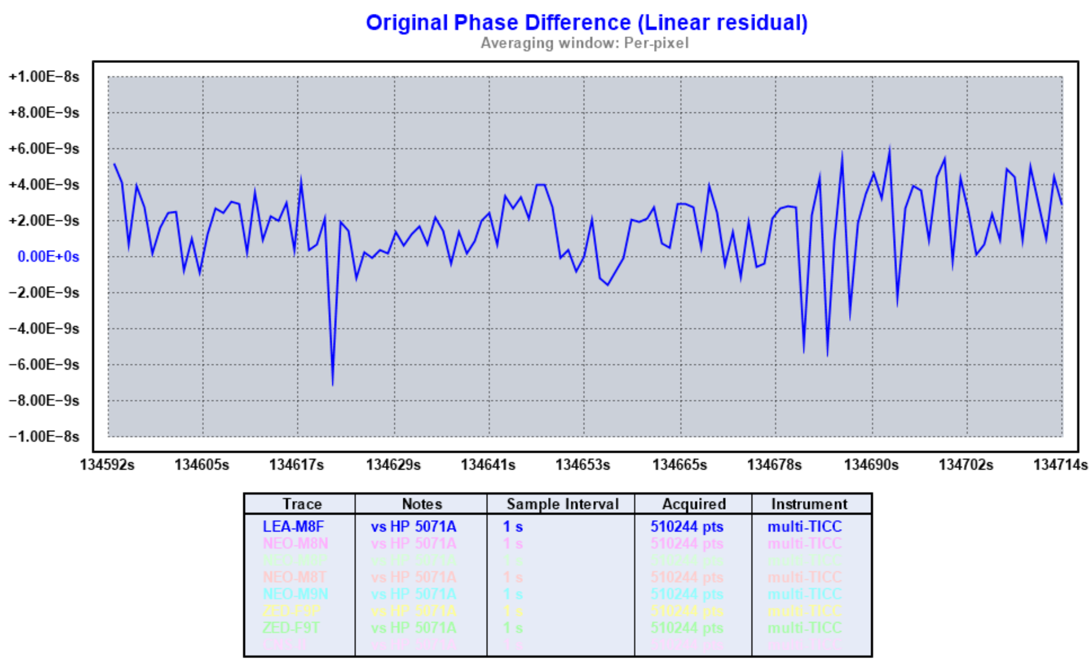

Figure 19: Zoomed Phase of LEA-M8F Receiver

Figure 17 and 18 show ADEV and phase performance for the LEA-M8F receiver, and Figure 19 shows a zoomed segment of phase data. The ADEV performance of the LEA-M8F is quite similar to the raw TIMEPULSE performance of the ZED-F9 series. Notable in the LEA-M8F ADEV plot is a slight ripple that is not present in the other receivers. Also note the occasional excursions of about 8 nanoseconds apparent in Figure 19. These seem to be the main limiting factor in the unit's ADEV performance.

2.5.2 LEA-M8F Frequency Performance

Characterizing the LEA-M8F's 30.72 MHz steered oscillator output requires different techniques than those used for a PPS signal. A Miles Labs TimePod 5330A phase noise test set[21], was used to characterize the Allan Deviation and phase noise of the M8F compared to the same 5071A Cesium frequency reference used for the other measurements in this paper. Using the same TimeLab software as before, it is also possible to show frequency as well as phase data, though it is not possible to zoom sufficiently to view cycle-by-cycle performance at this input frequency. The measurement system also allows plotting of RF phase noise data.

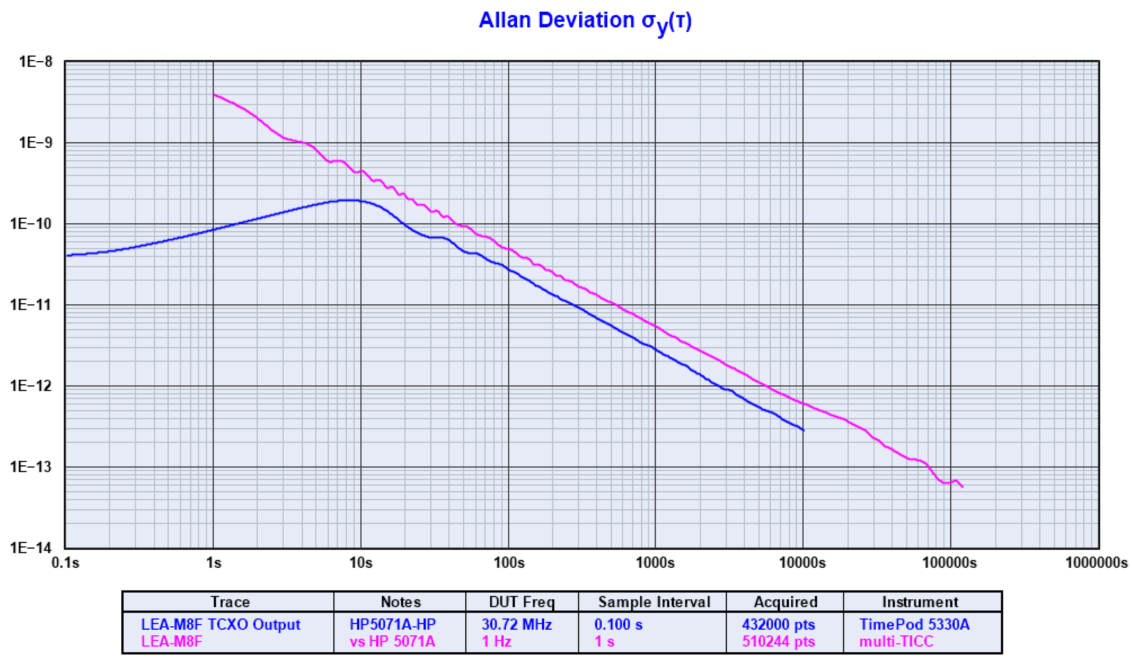

Figure 20: Allan Deviation of LEA-M8F Receiver

21 This unit is equivalent to the Micro-Semi 3120A (https://www.microsemi.com/product-directory/phase-noise-and-allen-deviation-testers/4131-3120a).

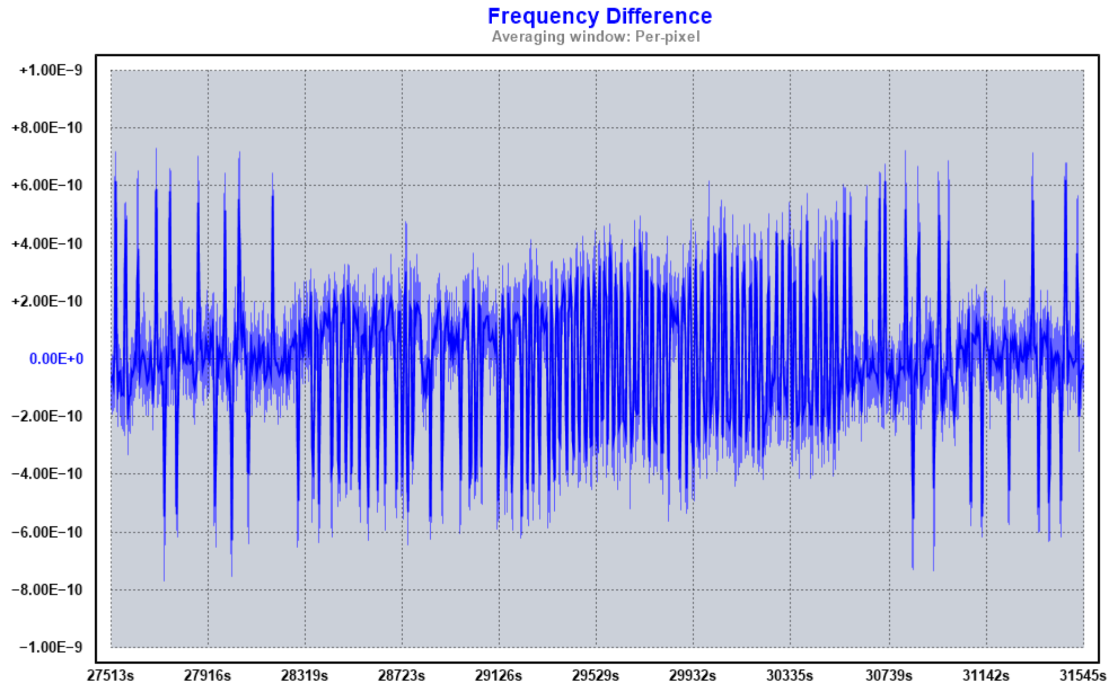

Trace	Notes	Input Freq	Input Amplitude	Duration	Elapsed	Instrument
u-Blox M8F default config (Unsaved)	HP5071A-HP	30.72 MHz	6.1 dBm	12h	12h	TimePod 5330A

Figure 21: LEA-M8F Relative Frequency vs. Cesium

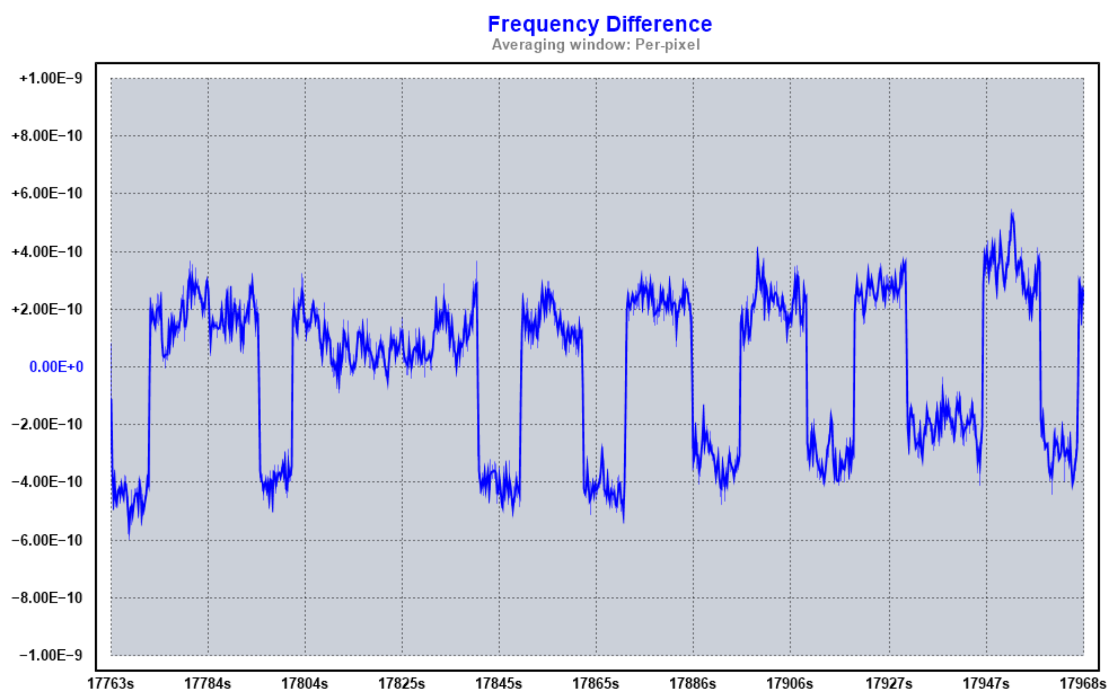

Trace	Notes	Input Freq	Input Amplitude	Duration	Elapsed	Instrument
u-Blox M8F default config (Unsaved)	HP5071A-HP	30.72 MHz	6.1 dBm	12h	12h	TimePod 5330A

Figure 22: Zoomed LEA-M8F Relative Frequency vs. Cesium

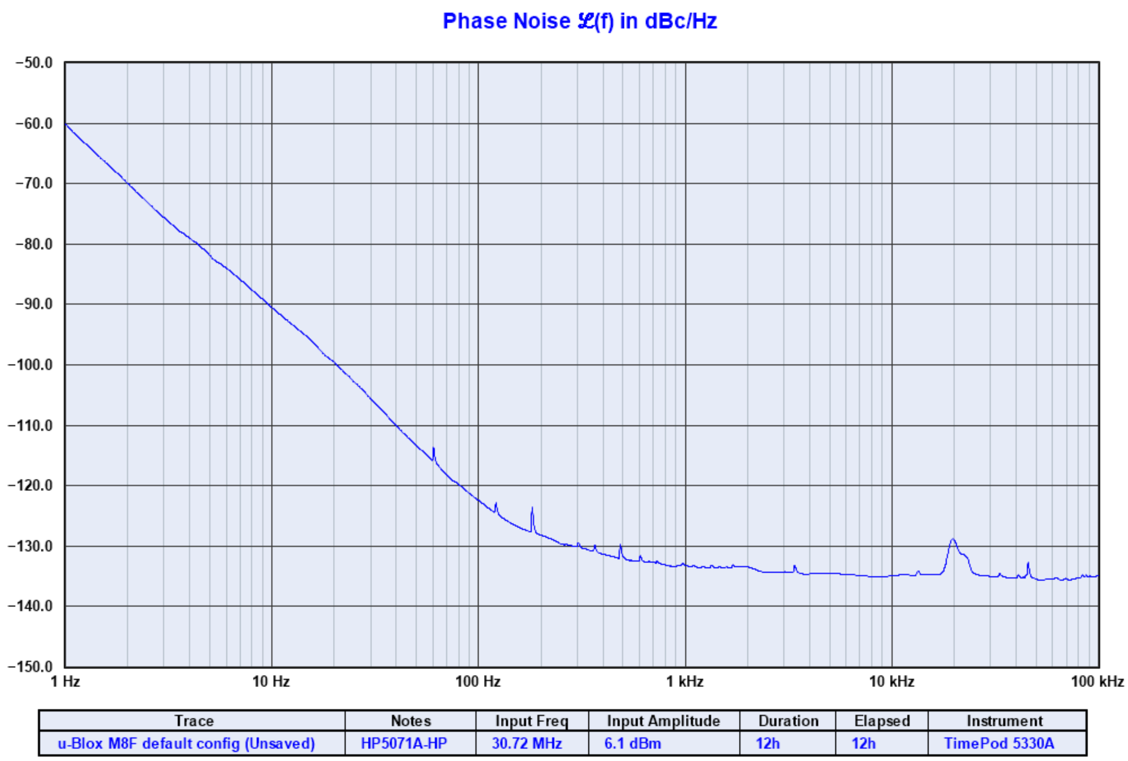

Figure 23: LEA-M8F Phase Noise

The blue line in Figure 20 plots the ADEV of the LEA-M8F steered oscillator output, while the violet line shows the 1 PPS TIMEPULSE output. The 30.72 MHz ADEV is significantly better than the TIMEPULSE output at short tau (<10 seconds) and slightly better at longer tau.

Figure 21 plots relative frequency difference of the LEA-M8F compared to the reference frequency standard over several thousand seconds. Note that the frequency record shows significant variability with numerous spikes of around $6x10^{-10}$ at varying intervals as well as a series of smaller excursions.

Figure 22 shows an expanded view of the frequency difference data from the LEA-M8F. This reveals that the output dithers approximately +/- 3 x 10^{-10} around the nominal frequency over periods of about 10 to 40 seconds. It thus appears that a form of pulse width modulation is used to obtain a nominal frequency that is “correct” on average. In fact, however, at any instant the frequency is either plus or minus about 0.3 parts per billion (“PPB”) from nominal. Whether this level of frequency jitter is acceptable in an RF application will depend upon the application. The frequency of this modulation indicates that the steering control loop time constant is quite short; its exact parameters are not known.

Figure 23 shows the phase noise of the 30.72 MHz oscillator from the LEA-M8F. The signal is relatively spur-free, but may be of marginal quality for high-performance RF applications, where a phase noise floor below -130 dBc/Hz is usually desirable.

Conclusion: The LEA-M8F is a unique design that might be useful in some situations where a direct RF signal output of reasonable spectral quality is desired. However, it does not yield results as good as those obtained by traditional “GPSDO” designs, and lack of ability to optimize control parameters limits its usefulness with external oscillators. The dithering used to obtain fine frequency control may be problematic in some applications.

3. SAWTOOTH OR QUANTIZATION ERROR CORRECTION

When a GPS receiver calculates its navigation solution, it can determine with great precision when the top of the second will occur. That knowledge is used to trigger the hardware PPS signal. But the hardware signal must come from somewhere, and it is usually derived from, and is synchronous with, the crystal oscillator that runs the receiver. If the clock runs at, for example, 25 MHz that means there is a timing pulse available every 40 nanoseconds and the receiver can provide an output pulse that is within 20 nanoseconds either side (*i.e.*, early or late) of the calculated time marker.[22] If the clock frequency is increased, the PPS granularity is reduced and the time pulse can have even smaller ambiguity. A higher clock rate is the probable explanation for the ZED-F9 series receivers' smaller jitter and better short-term ADEV than the M8 series.

But no clock is ever exactly on frequency, and the receiver has to compensate for the difference between the hardware clock and the GPS constellation clock. It does so by adding or dropping hardware clock ticks between subsequent PPS output pulses to keep the hardware signal aligned as closely as possible to the software-determined second marker. In other words, if the clock runs at 25 MHz instead of counting 25 million ticks between PPS outputs, the receiver might count 25,000,001 or 24,999,999. The result of these changes causes an effect that looks in short-term plots of GPS phase like a sawtooth pattern. u-blox calls this effect "quantization error" and it is visible in the sawtooth on the left and right sides of Figure 24.

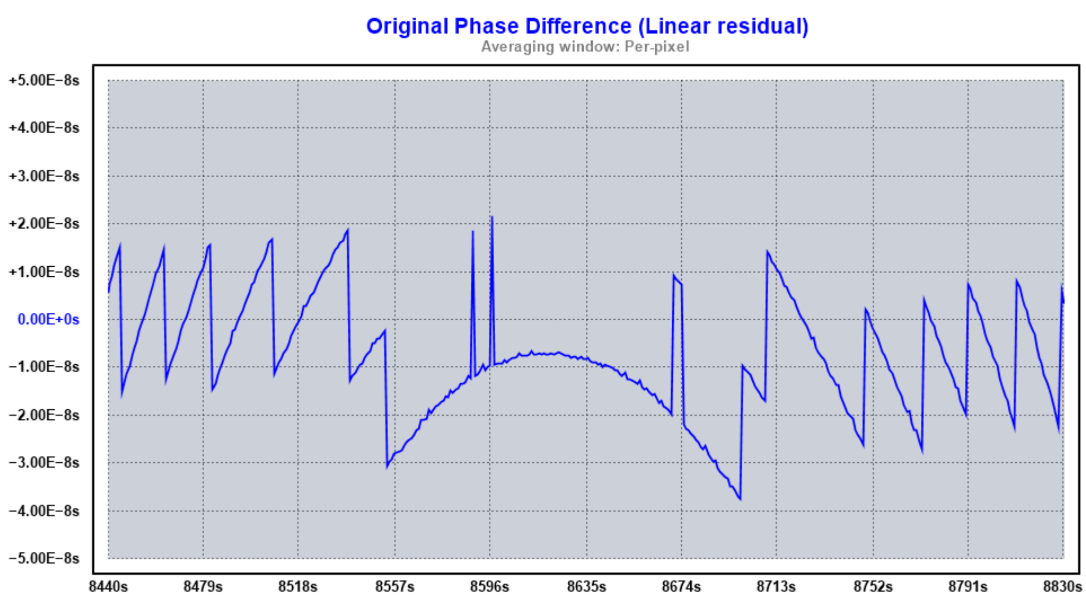

Trace	Notes	Sample Interval	Acquired	Instrument
HP5065A vs. GPS PPS	vis CNS-II and TICC	1 s	13770 pts	

Figure 24: GPS PPS "Hanging Bridge"

22 That is theoretical accuracy; other factors may add fixed or variable offsets to the actual output pulse timing.

This jitter adds noise and worsens the ADEV of the PPS output, especially at short tau. If the frequency offset remains constant the jitter pattern will remain a sawtooth, and increased averaging can reduce or eliminate its effects so it does not impair long-term ADEV.

Over time, though, the oscillator will drift and have other instabilities such as those caused by temperature changes or frequency drift, and those changes can cause the phase relationship between the hardware signal and the computed time mark to change. A slow drift can result in the sawtooth turning into interesting patterns such as the “hanging bridge” shown in the middle of Figure 24. Patterns such as these result in short-term offsets above or below the mean PPS frequency, which reduce the overall quality of the timing result and worsen ADEV.

These effects can be mitigated in several ways. Higher quality internal oscillators (for example, temperature-controlled crystal oscillators, or “TCXOs”) can slow down frequency changes, making the sawtooth cycle longer (which may or may not be a good thing, if the result is to spend more time in a “hanging bridge” condition), and using higher internal frequencies improves resolution. For example, using a 100 MHz clock reduces the granularity to 10 nanoseconds.

While the sawtooth effect is very difficult to eliminate in hardware, software can be used to mitigate it. Many modern GPS units make available a message in the output data stream which predicts the offset of the upcoming hardware pulse from the GPS top-of-second. In the u-blox receivers, this is called a “quantization error” or “qErr” message and it is very effective to reduce the effect of hardware errors in cases where it is feasible to apply a software correction to the hardware PPS results.

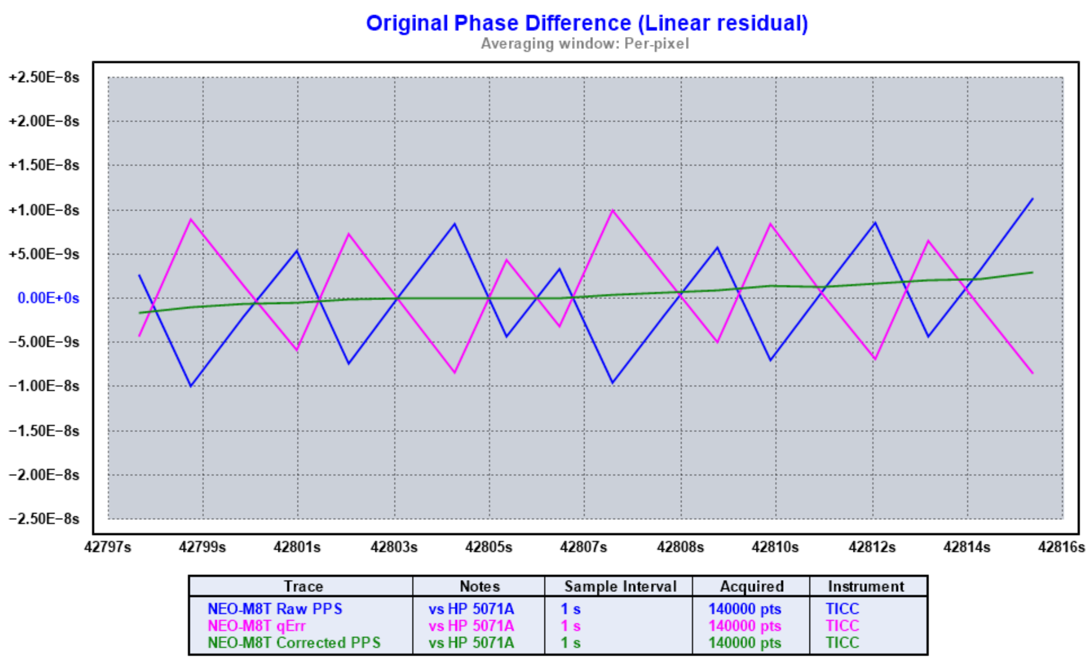

Figure 25: Quantization Error and Correction

Figure 25 zooms in on the phase plot for an M8T with two traces added, one showing the value of the quantization error (“qErr”) value received via the receiver serial port (in pink), and the second showing the result of subtracting the qErr value from the phase value (in green). The qErr data almost exactly offsets the raw PPS jitter. Subtracting the two yields the green line, in which almost all jitter has been removed, showing the effectiveness of this technique (and making visible via the upward slope a small offset between GPS and local reference at that point in time).

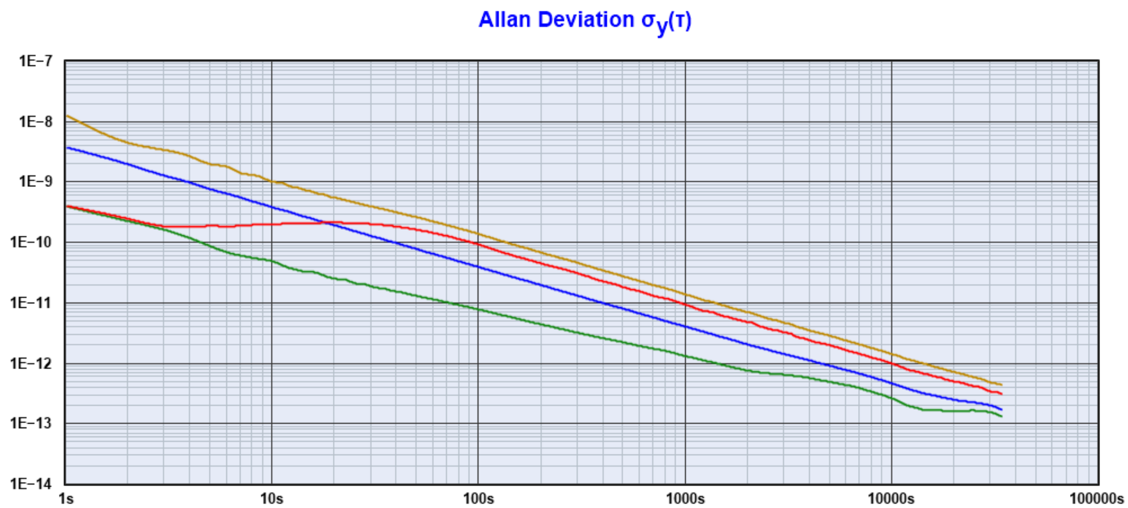

Trace	Notes	Sample Interval	Acquired	Instrument
ZED-F9T Raw PPS	vs HP 5071A	1 s	140000 pts	TICC
ZED-F9T qErr	vs HP 5071A	1 s	140000 pts	TICC
ZED-F9T Corrected PPS	vs HP 5071A	1 s	140000 pts	TICC
NEO-M8T Corrected PPS	vs HP 5071A	1 s	140000 pts	TICC
NEO-M8T qErr	vs HP 5071A	1 s	140000 pts	TICC
NEO-M8T Raw PPS	vs HP 5071A	1 s	140000 pts	TICC

Figure 26: Allan Deviation After qErr Correction

Figure 26 shows before-and-after ADEV performance of both the NEO-M8T and ZED-F9T when applying sawtooth correction. Note that the M8T shows the most dramatic improvement at short tau, and beyond 100 seconds a much more modest one. Conversely, the F9T shows significant improvement out to much longer tau. Overall, the dual-frequency receivers perform about an order of magnitude better than the single-frequency ones.

Looking more closely at the ADEV results, at 3 seconds and below, the two receivers yield almost identical corrected results. It is believed that within this regime the more predictable resolution-limited jitter dominates, and both receivers can correct for this very precisely. At longer tau, however, GPS noise as well as hardware limitations enter the picture, and because the F9T uses its dual-frequency capability to reduce GPS noise, the quantization error remains visible at longer taus, and the effects of the error correction also remain apparent.

Conclusion: Quantization error correction is a powerful tool to improve the short-term jitter performance of GPS receivers. It is at least as beneficial for dual-frequency designs as for single-frequency ones.

4. FIXED POSITION (“0-D”) VS. 3-D NAVIGATION TIMING RESULTS

One difference between standard GPS receivers and units specialized for timing purposes is the ability to enter into a “timing” or “0D” navigation mode, where the receiver’s location is fixed and only the time is calculated in the solution. The receiver location can be determined either autonomously via a site-survey function that averages position readings over a period of time, or by manual entry of the known location into the receiver parameters.

The 0D navigation mode should produce better timing results because there is only one variable being solved for, rather than four (three physical dimensions plus time) as in a normal position fix. How much difference in performance does it actually make?

As shown in Figures 27 through 29, using the 0D mode does improve the Allan Deviation of the TIMEPULSE output for tau greater than about 30 seconds, though there is little difference in the short term. The dual frequency ZED-F9T receiver shows much greater benefit than the single-frequency NEO-M8T. The phase plots show that 0D mode subjectively wanders less over longer time intervals.

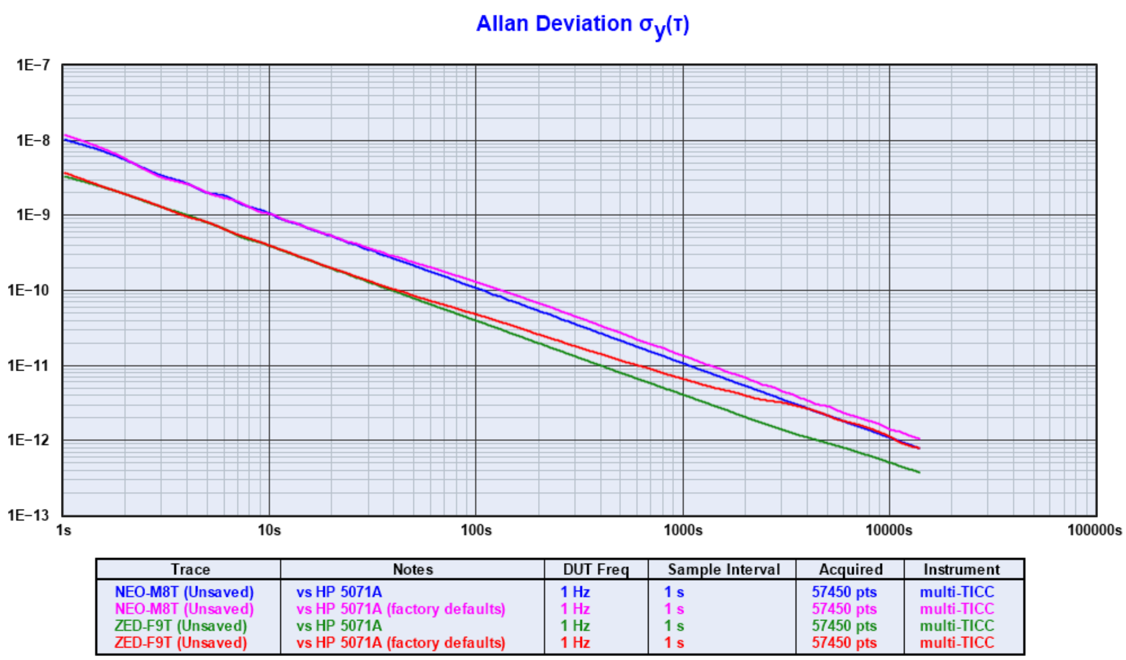

Figure 27: Allan Deviation of Timing vs. 3D Navigation Solution

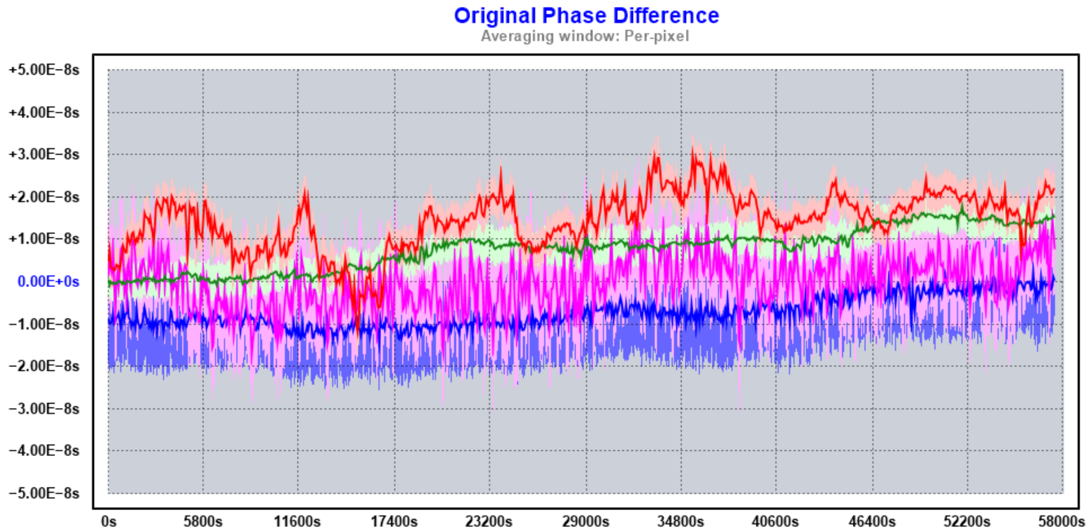

Trace	Notes	DUT Freq	Sample Interval	Acquired	Instrument
NEO-M8T (Unsaved)	vs HP 5071A	1 Hz	1 s	57450 pts	multi-TICC
NEO-M8T (Unsaved)	vs HP 5071A (factory defaults)	1 Hz	1 s	57450 pts	multi-TICC
ZED-F9T (Unsaved)	vs HP 5071A	1 Hz	1 s	57450 pts	multi-TICC
ZED-F9T (Unsaved)	vs HP 5071A (factory defaults)	1 Hz	1 s	57450 pts	multi-TICC

Figure 28: Raw Phase of Timing vs. 3D Navigation Solution

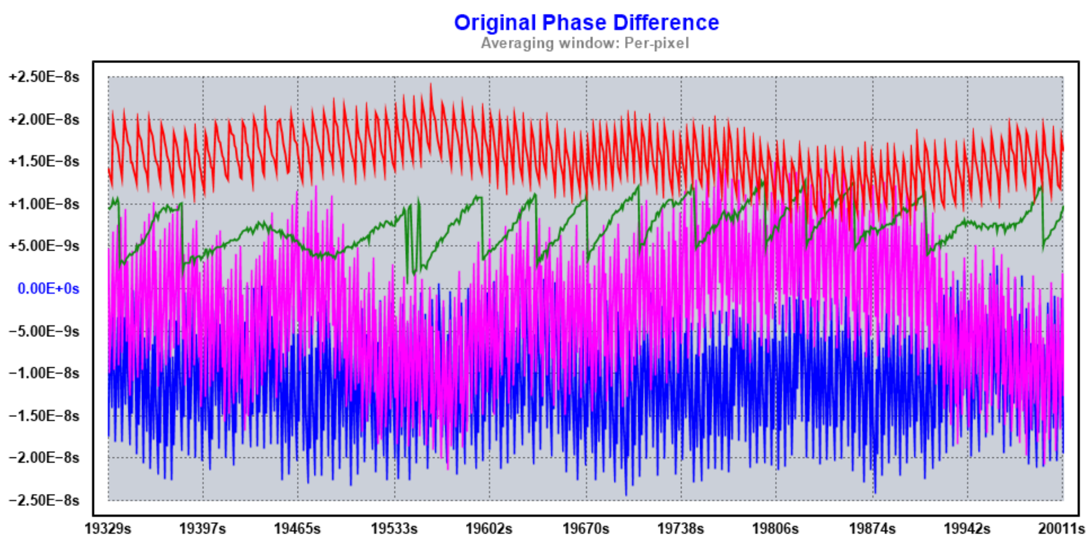

Trace	Notes	DUT Freq	Sample Interval	Acquired	Instrument
NEO-M8T (Unsaved)	vs HP 5071A	1 Hz	1 s	57450 pts	multi-TICC
NEO-M8T (Unsaved)	vs HP 5071A (factory defaults)	1 Hz	1 s	57450 pts	multi-TICC
ZED-F9T (Unsaved)	vs HP 5071A	1 Hz	1 s	57450 pts	multi-TICC
ZED-F9T (Unsaved)	vs HP 5071A (factory defaults)	1 Hz	1 s	57450 pts	multi-TICC

Figure 29: Zoomed Phase of Timing vs. 3D Navigation Solution

Conclusion: Use of “timing” or “0D” navigation solution mode provides improved timing performance, particularly over longer measurement intervals.

5. RESULTS AT HIGHER TIMEPULSE RATES

The u-blox receivers compared here all allow the TIMEPULSE output to be set to much higher rates than one pulse per second, typically 10 MHz or greater. Some receivers also allow a higher navigation solution rate than one per second (8 Hz for the NEO-M9N and 25 Hz for the dual-frequency ZED-F9T). Accordingly, the inexpensive NEO-M8N, its replacement, the NEO-M9N, and high-end ZED-F9T were tested to determine how they perform at a 10 MHz TIMEPULSE rate, and also at higher navigation solution rates where supported. The test configuration was the same as that used for the LEA-M8F measurements reported in Section 2.7.

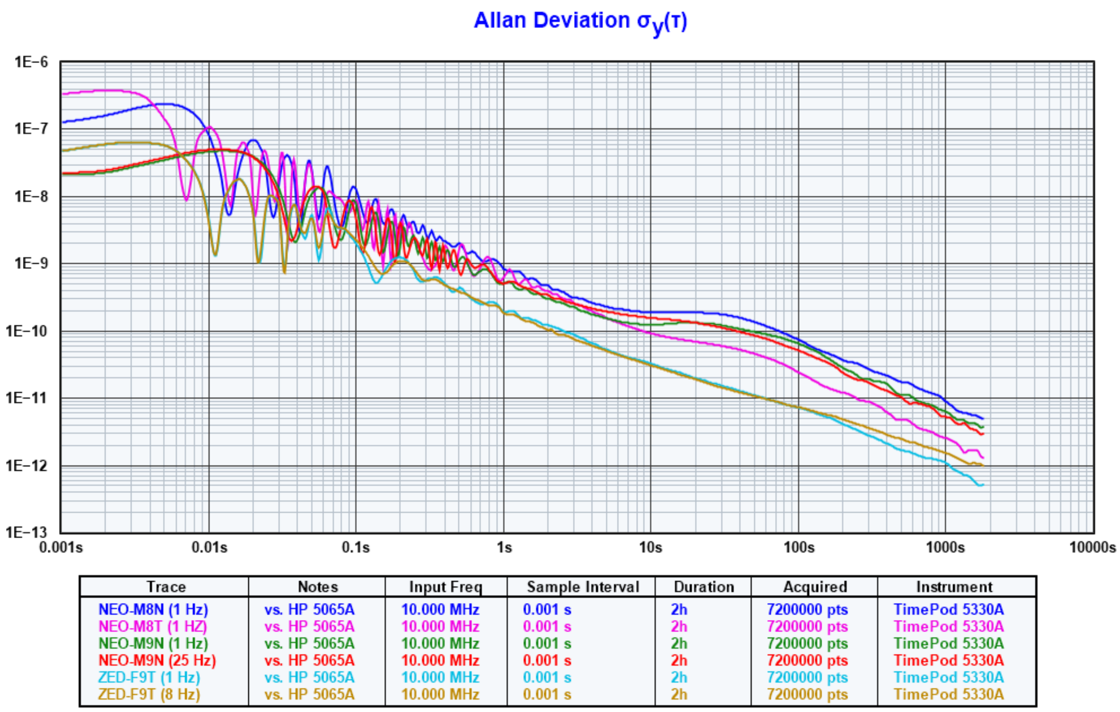

Figure 30: Allan Deviation at 10 MHz TIMEPULSE Rate

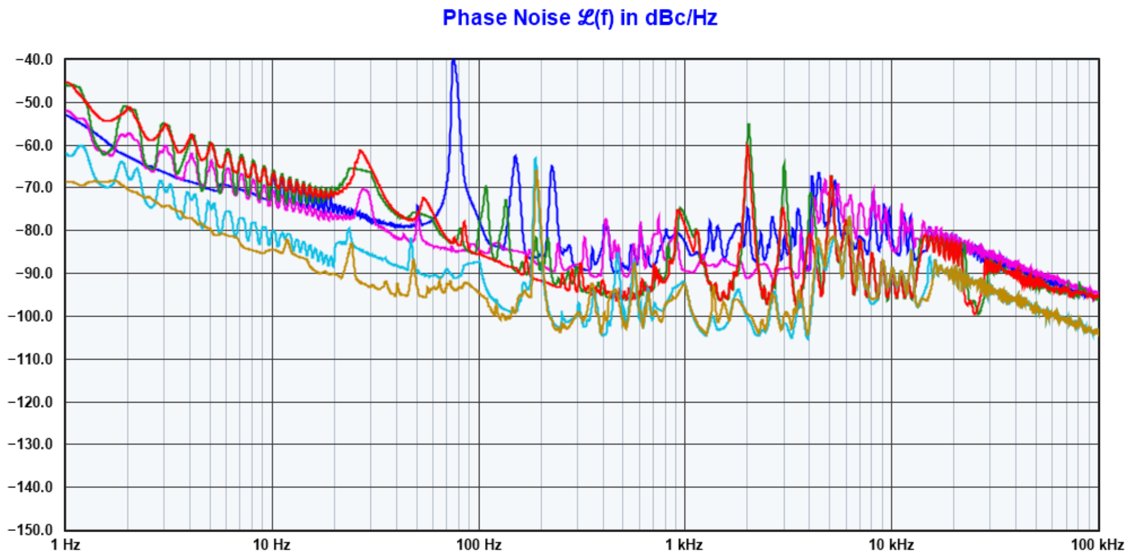

Trace	Notes	Input Freq	Sample Interval	Duration	Acquired	Instrument
NEO-M8N (1 Hz)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A
NEO-M8T (1 HZ)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A
NEO-M9N (1 Hz)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A
NEO-M9N (25 Hz)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A
ZED-F9T (1 Hz)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A
ZED-F9T (8 Hz)	vs. HP 5065A	10.000 MHz	0.001 s	2h	7200000 pts	TimePod 5330A

Figure 31: Phase Noise at 10 MHz TIMEPULSE Rate

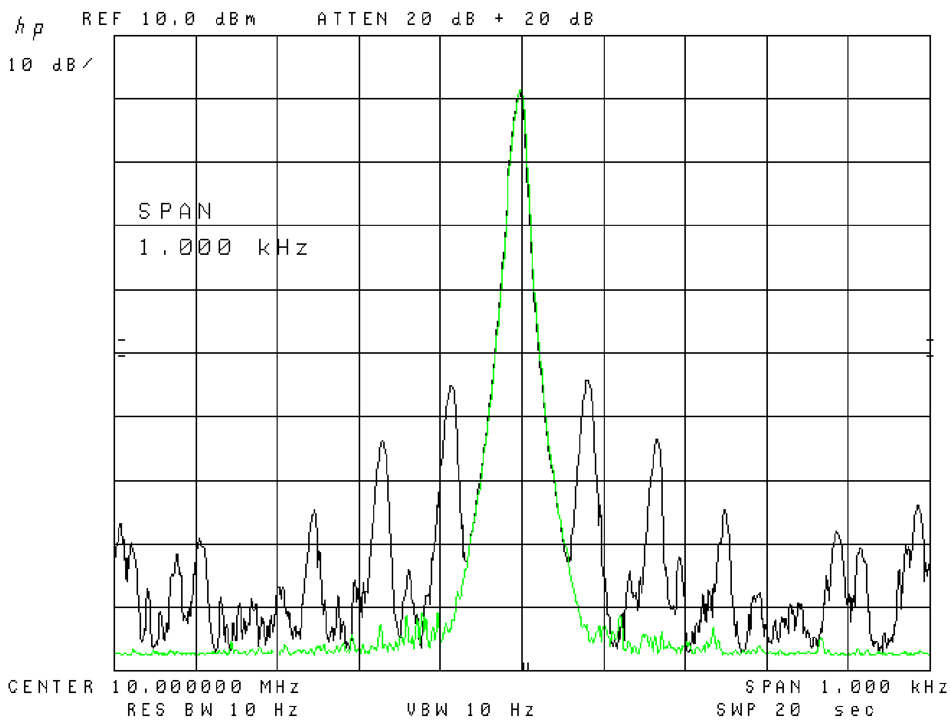

Figure 32: Spectrum Analyzer View of ZED-F9T (Black) and HP 8642A (Green) at 10 MHz, 1 kHz Span

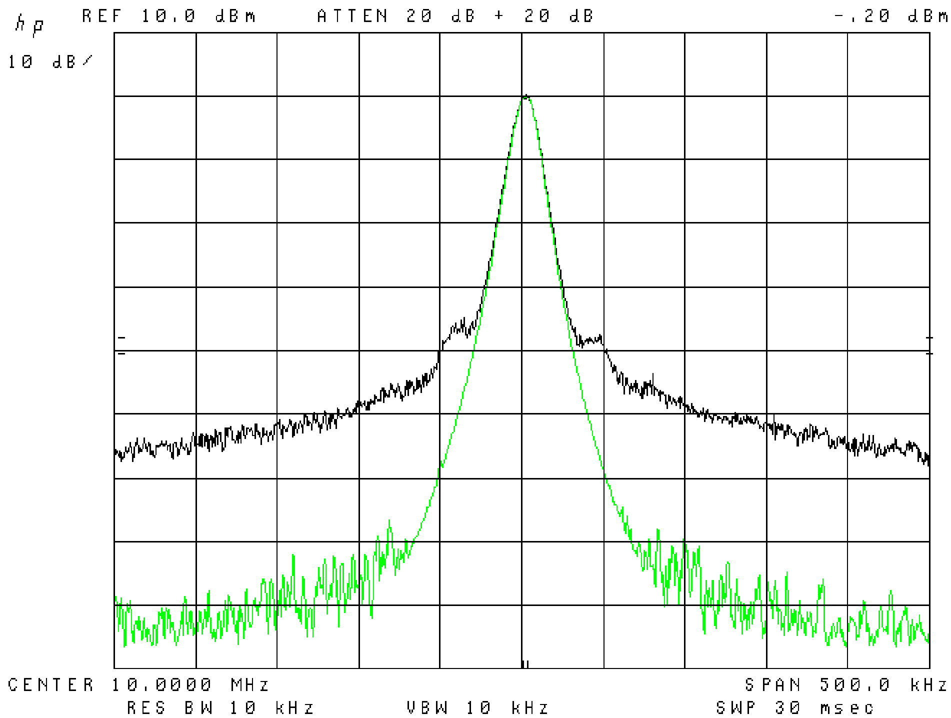

Figure 33: Spectrum Analyzer View of ZED-F9T (Black) and HP 8642A (Green) at 10 MHz, 500 kHz Span

Figure 30 shows the ADEV of the receivers when set to 10 MHz output and 1 Hz navigation solution rates, as well as results for the higher rates supported in the 9-series units.[23] In general, the ADEV of a 10 MHz TIMEPULSE is somewhat better than at one pulse per second. However, the "ringing" oscillations at tau less than 1 second are warnings that not all is well, particularly since a signal at 10 MHz is likely to be used in an RF system where signal purity and spur-free performance is important.

A phase-noise view of the data, as shown in Figure 31, reveals that there is indeed cause for concern. In particular, not only is the performance at close offsets substandard compared to most RF signal sources, but the noise floor is very high at about -105 dBc/Hz at 100 kHz offset. These signals are clearly ***not*** suitable for use in most RF applications. Figures 32 and 33 show screen captures from an HP 8568B spectrum analyzer comparing the ZED-F9T 10 MHz TIMEPULSE output (white) with a laboratory signal generator (violet). Figure 31 is a close-up showing a 1 kHz span, while Figure 32 shows 500 kHz.

Conclusion: While it would be unreasonable to expect that any GPS receiver signal evaluated this way to be as clean as a laboratory-grade generator, the TIMEPULSE noise is 30 to 40 dB higher and is probably unusable for RF applications.

23 Note that the X axis begins at 0.001 second in this plot, where the previous plots have all begun at 1 second. This is because the most interesting results appear in this short-tau regime, where frequency stability and phase noise measurements start to cross over.

6. THE "EXTERNAL INTERRUPT" INPUT

Several of the u-blox receivers have one or more "EXT INT" (external interrupt) inputs that can be used to timestamp an external signal against GPS time. This input was tested using a 1 PPS signal derived from the 5071A Cesium frequency standard, logging the output of the u-blox "UBX-TIM-TM2" binary message to disk for later processing.

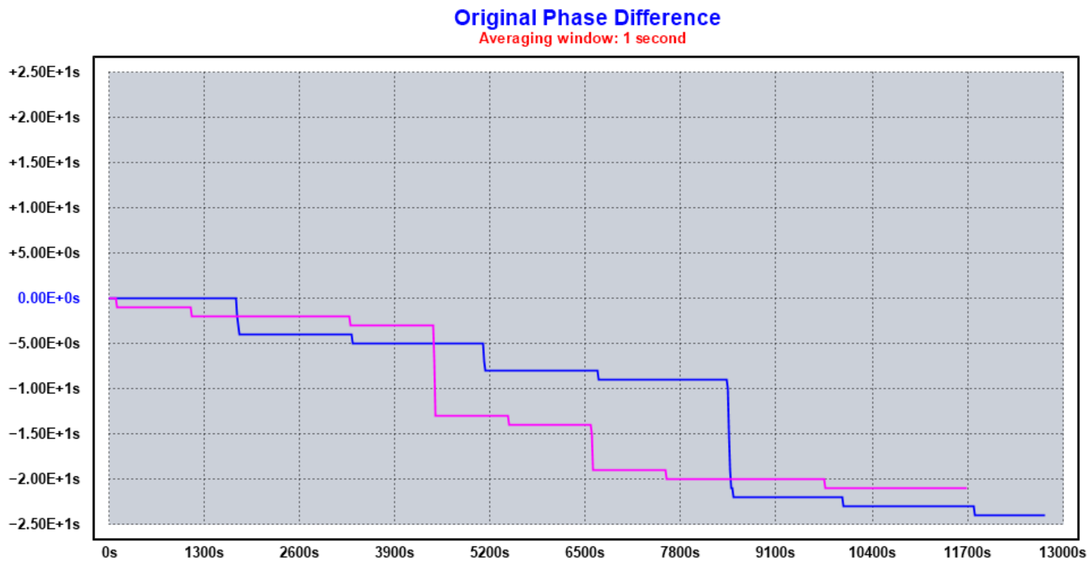

Trace	Notes	Sample Interval	Acquired	Instrument
NEO-M8T EXT INT	vs HP 5071A	1 s	12754 pts	TICC
ZED-F9P EXT INT	vs HP 5071A	1 s	11684 pts	TICC

Figure 34: Raw Phase of EXT INT Inputs

Figure 34 shows the phase records of the NEO-M8T (blue) and ZED-F9P (magenta)[24] with the EXTINT feature activated and driven by the PPS signal from an HP 5071A Cesium standard. There is clearly a problem here; the stairsteps in the phase record show jumps of several seconds.

It appears that the receivers in this mode sometimes both duplicate and miss input pulses. Here is a portion of the data logged from the UBX-TIM-TM2 message:

```
331370.000000418 10 456
331371.000000420 10 457
331371.000000420 10 457          # duplicate
                                 # 33172 missing
331373.000000418 10 459
331373.000000418 10 459          # duplicate
331374.000000420 10 460
```

24 It would have been more consistent to use the ZED-F9T timing receiver for this test, but the evaluation board for that unit does not make the EXT INT input accessible.

These repeated and missing data points appear with considerable frequency in both the NEO-M8T and ZED-F9P results. If there were only duplicates, it would be easy to filter them out, but the missing data points are very difficult to handle and ultimately make it impossible to create a complete phase record.

Figures 35 through 37 show results after selecting a data segment in which none of these glitches occurred. Using the glitch-free data, the results look broadly similar to those of the TIMEPULSE output, which is what would be expected.

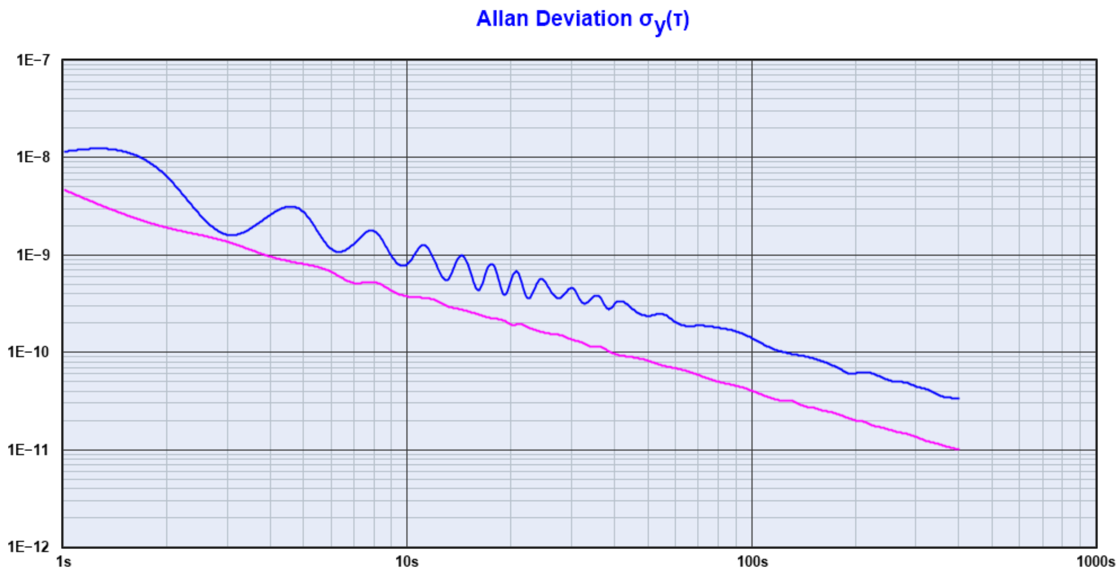

Trace	Notes	Sample Interval	Acquired	Instrument
NEO-M8T EXT INT (trimmed)	vs HP 5071A	1 s	1791 pts	TICC
ZED-F9P EXT INT (trimmed)	vs HP 5071A	1 s	1791 pts	TICC

Figure 35: Allan Deviation of EXT INT In Glitch-Free Segment

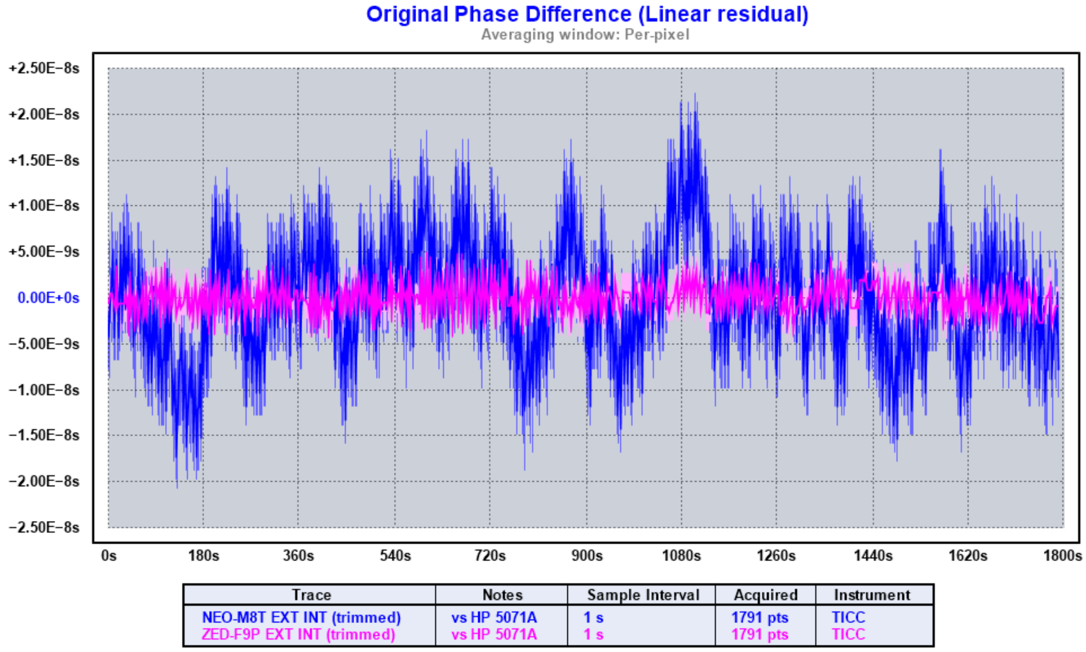

Figure 36: Raw Phase of EXT INT In Glitch-Free Segment

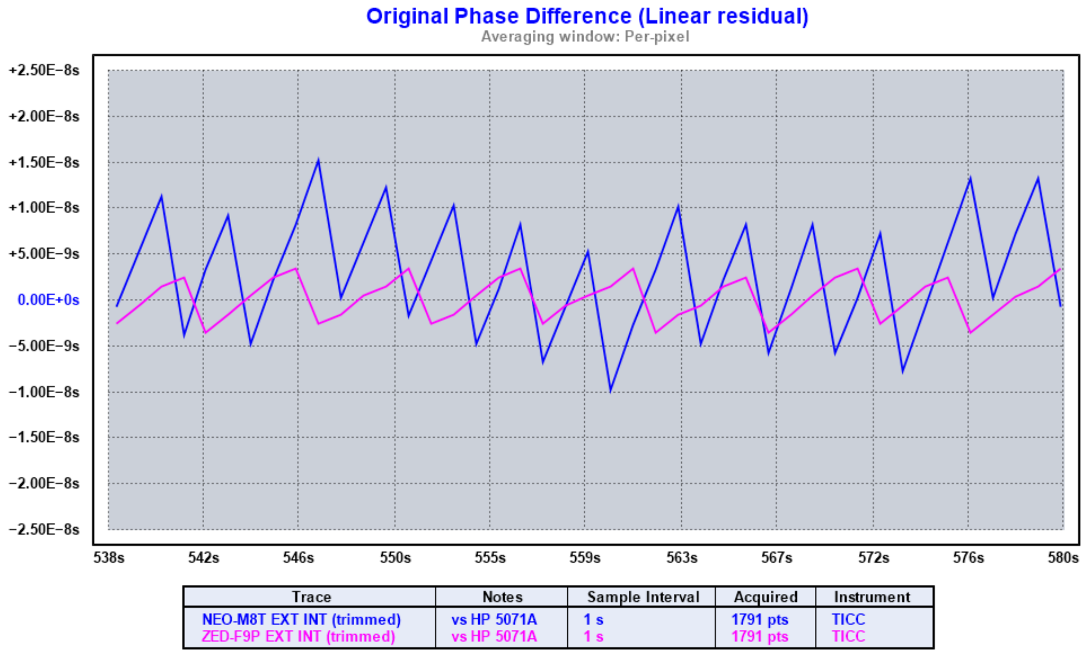

Figure 37: Zoomed Raw Phase of EXT INT In Glitch-Free Segment

Conclusion: If it were not for these issues, the performance of the EXT INT measurement appears similar to that of the TIMEPULSE output and could be useful for timestamping, as shown by figures 35 through 37. Investigation into the causes of the glitches continues.

7. IMPLICATIONS FOR GPS DISCIPLINED OSCILLATORS

A GPS Disciplined Oscillator ("GPSDO") uses the pulse-per-second output of a GPS receiver to steer the frequency of a crystal oscillator. As shown in this paper, the GPS PPS signal has noise limiting its short-term performance, but over the long term that signal tracks the ensemble of atomic clocks used to control the GPS satellite constellation.

A free-running crystal oscillator can have very low noise in the short term, but it needs to be adjusted to its nominal frequency, and suffers from temperature sensitivity and long-term frequency drift due to aging, both of which limit its ultimate accuracy and stability.

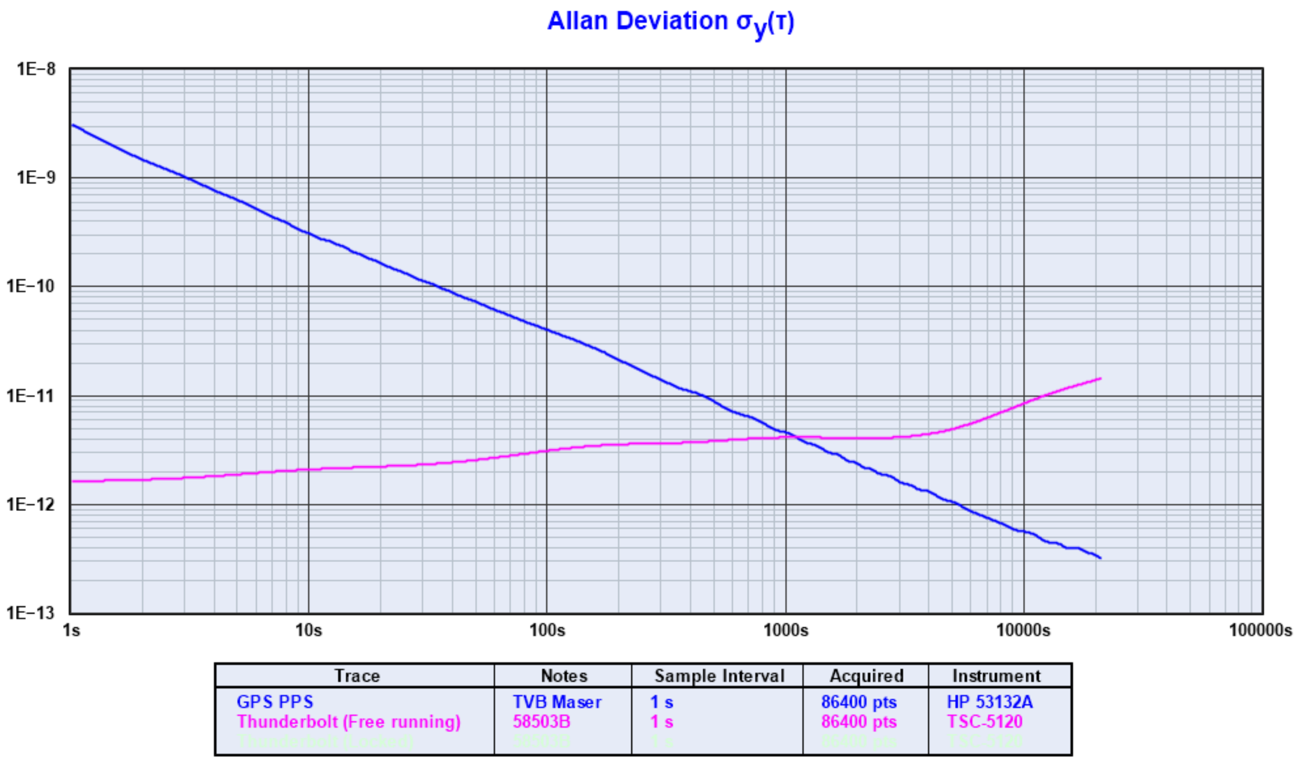

Figure 38: Allan Deviation Intercept of GPS and OCXO[25]

Figure 38 shows how the performance of a crystal oscillator is better at short measurement intervals, while the GPS is more stable over longer tau. The purpose of the GPSDO is to obtain the best of both worlds by steering the crystal oscillator frequency to follow the GPS PPS frequency. In effect, the oscillator acts as a "flywheel" to smooth the GPS noise.

This is normally accomplished by dividing the crystal oscillator output to create a PPS signal, comparing the timing of that signal to the GPS PPS, and using a phase lock loop to steer the crystal so that its PPS signal tracks that of the GPS receiver. The bandwidth, or time constant, of the loop controls where each input signal is dominant in the output, and the goal is to set the bandwidth so that the transition occurs at the intercept point to maintain the lowest ADEV

25 Data courtesy of Tom Van Baak

over the entire range of tau. With a high quality oscillator and GPS as shown in Figure 38, the time constant is typically in the range of several hundred to two thousand or more seconds.

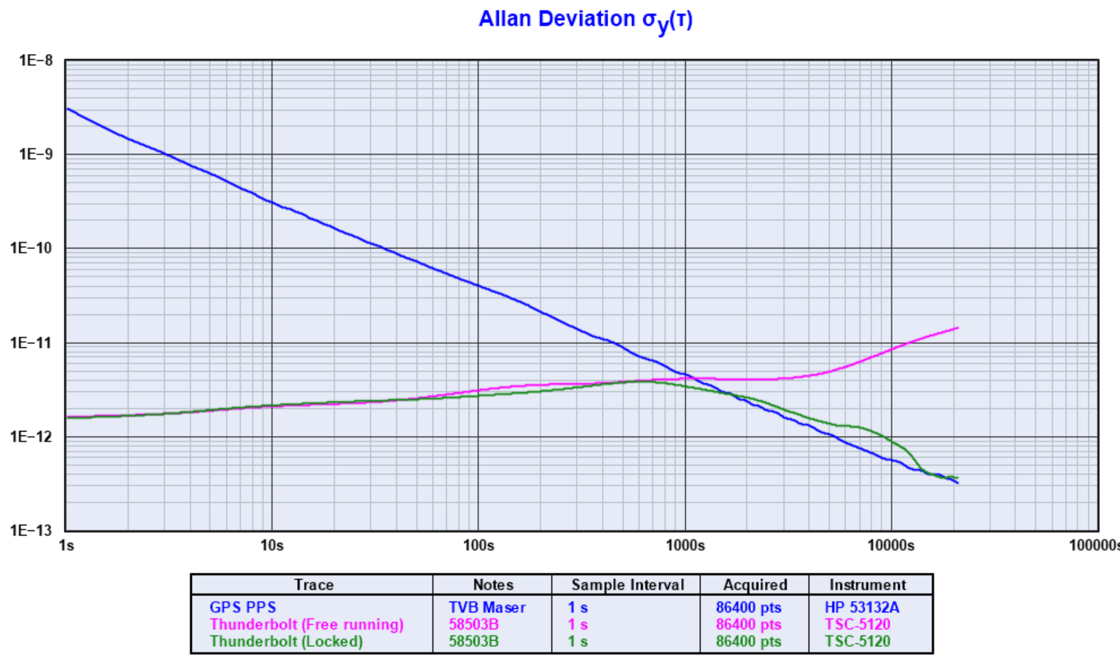

Figure 39: GPSDO Realization of Intercept Point[26]

The green trace in Figure 39 shows the actual performance of a GPSDO overlaid on plots of the free-running oscillator, and GPS PPS, performance. This high-quality GPSDO shows the results obtainable when the underlying signals sources are good, and the control loop is properly tuned.

If the GPS PPS noise is lowered, the loop time constant can be decreased, causing the GPS to dominate at shorter tau. This means that the long term performance of the crystal oscillator becomes less important, and may allow use of a lower cost oscillator so long as its short-term stability is acceptable for the application.

26 Data courtesy of Tom Van Baak

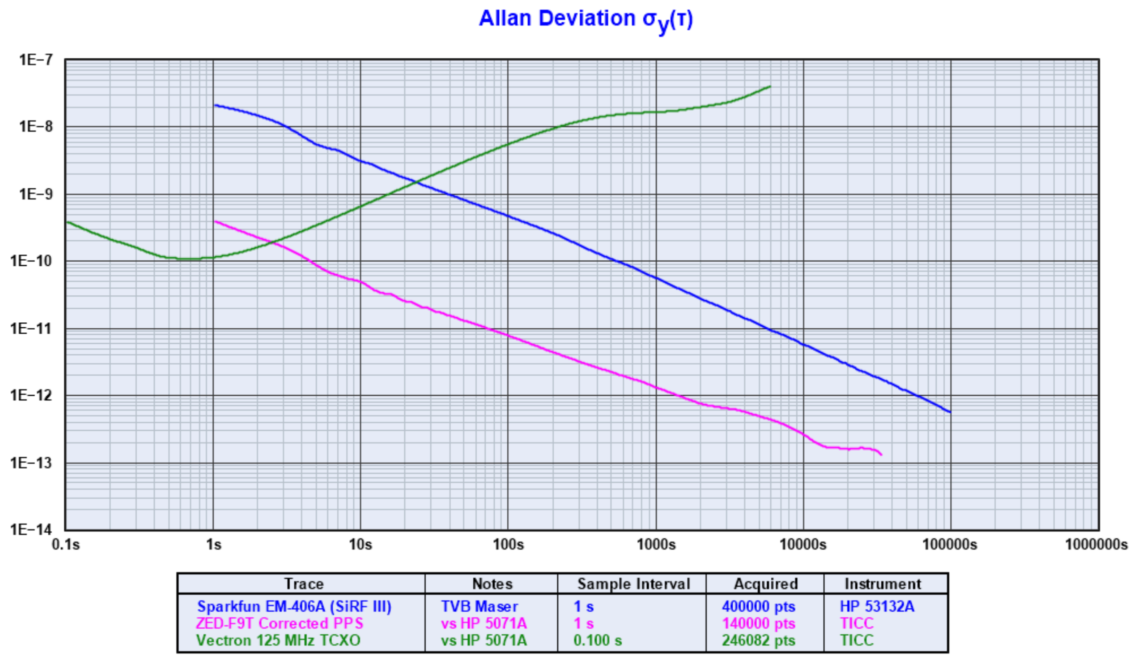

Figure 40: Allan Intercept of Low-Performance Oscillator

Figure 40 shows a less stable, less expensive, oscillator,[27] and also shows the impact of using a lower noise PPS source for control. With a mediocre GPS, the intercept point occurs at about 25 seconds, and the worst-case Allan Deviation is about 1.5×10^{-9}. By using the ZED-F9T receiver and applying sawtooth correction, the intercept point occurs at less than three seconds, and the worst-case Allan Deviation is about 2×10^{-10} – nearly an order of magnitude better.

Conclusion: Use of a dual-frequency GPS receiver together with application of sawtooth error correction may allow a GPS to be implemented with a lower-cost oscillator.

27 This is a VHF temperature controlled crystal oscillator (TCXO) running at 125 MHz which is optimized more for low phase noise than for frequency stability.

8. POSITIONING PERFORMANCE

Although the focus of this paper is on timing, it is appropriate to describe the positioning performance of these units. All provide both standard NMEA output data, as well as proprietary binary format messages. Some have a “high precision” NMEA output that provides seven decimal places of latitude and longitude precision compared to four places in the standard mode. Some have an RTK processing engine that allows real time corrections to be applied to the output solution. Most support SBAS (satellite-base augmentation system) corrections with no external hardware, though oddly the two “P” series positioning receivers do not. The “8” series receivers support concurrent GPS and GLONASS operation, while the “9” series receivers allow concurrent reception of four constellations (GPS, GLONASS, BeiDou and Galileo).

A full exploration of GPS positioning performance capabilities and tools could occupy many pages. The following is not intended to be a comprehensive report and should be viewed instead as a limited comparison of performance in selected measurements.

8.1 Autonomous Positioning

As an initial test, all seven units were connected to the same antenna[28] and set to their factory default condition, except that NMEA high precision (seven decimal place latitude and longitude resolution) was enabled where available and SBAS and QZSS augmentation were turned off. The 8-series receivers used the GPS and GLONASS constellations, while the 9-series receivers had all four constellations enabled.

The NMEA “GGA” sentence (“GPS Fix Data”) from each receiver was simultaneously logged to a disk file for 12 hours.[29] That data was then processed by the “gpsprof” program that is provided with the gpsd software suite[30] to generate a scatter plot and Circular Error Probability statistics for each receiver.

Circular Error Probable (“CEP”) is a generally accepted metric for GPS positioning performance.[31] It is roughly the radius of a circle containing a specified percentage of all data positions logged. 95 percent is most commonly used value. The gpsprof program also calculates a similar value for Elevation Probability (EP). Figure 41 shows the scatter plots and CEP and EP data for each receiver.

---

28 As noted earlier, the Trimble Zephyr Geodetic antenna used is capable of L1/L2 operation, but is not optimized for the carrier frequencies used by GLONASS and BeiDou. However, all the receivers were able to lock multiple GLONASS satellites with good carrier-to-noise ratios.

29 To obtain the best average position, it would be better to collect data for 24 or 48 hours, but time did not permit that long a measurement.

30 https://gpsd.gitlab.io/gpsd/index.html

31 See, *e.g.*, https://www.gpsworld.com/gps-accuracy-lies-damn-lies-and-statistics/; http://www.dtic.mil/dtic/tr/fulltext/u2/a199190.pdf

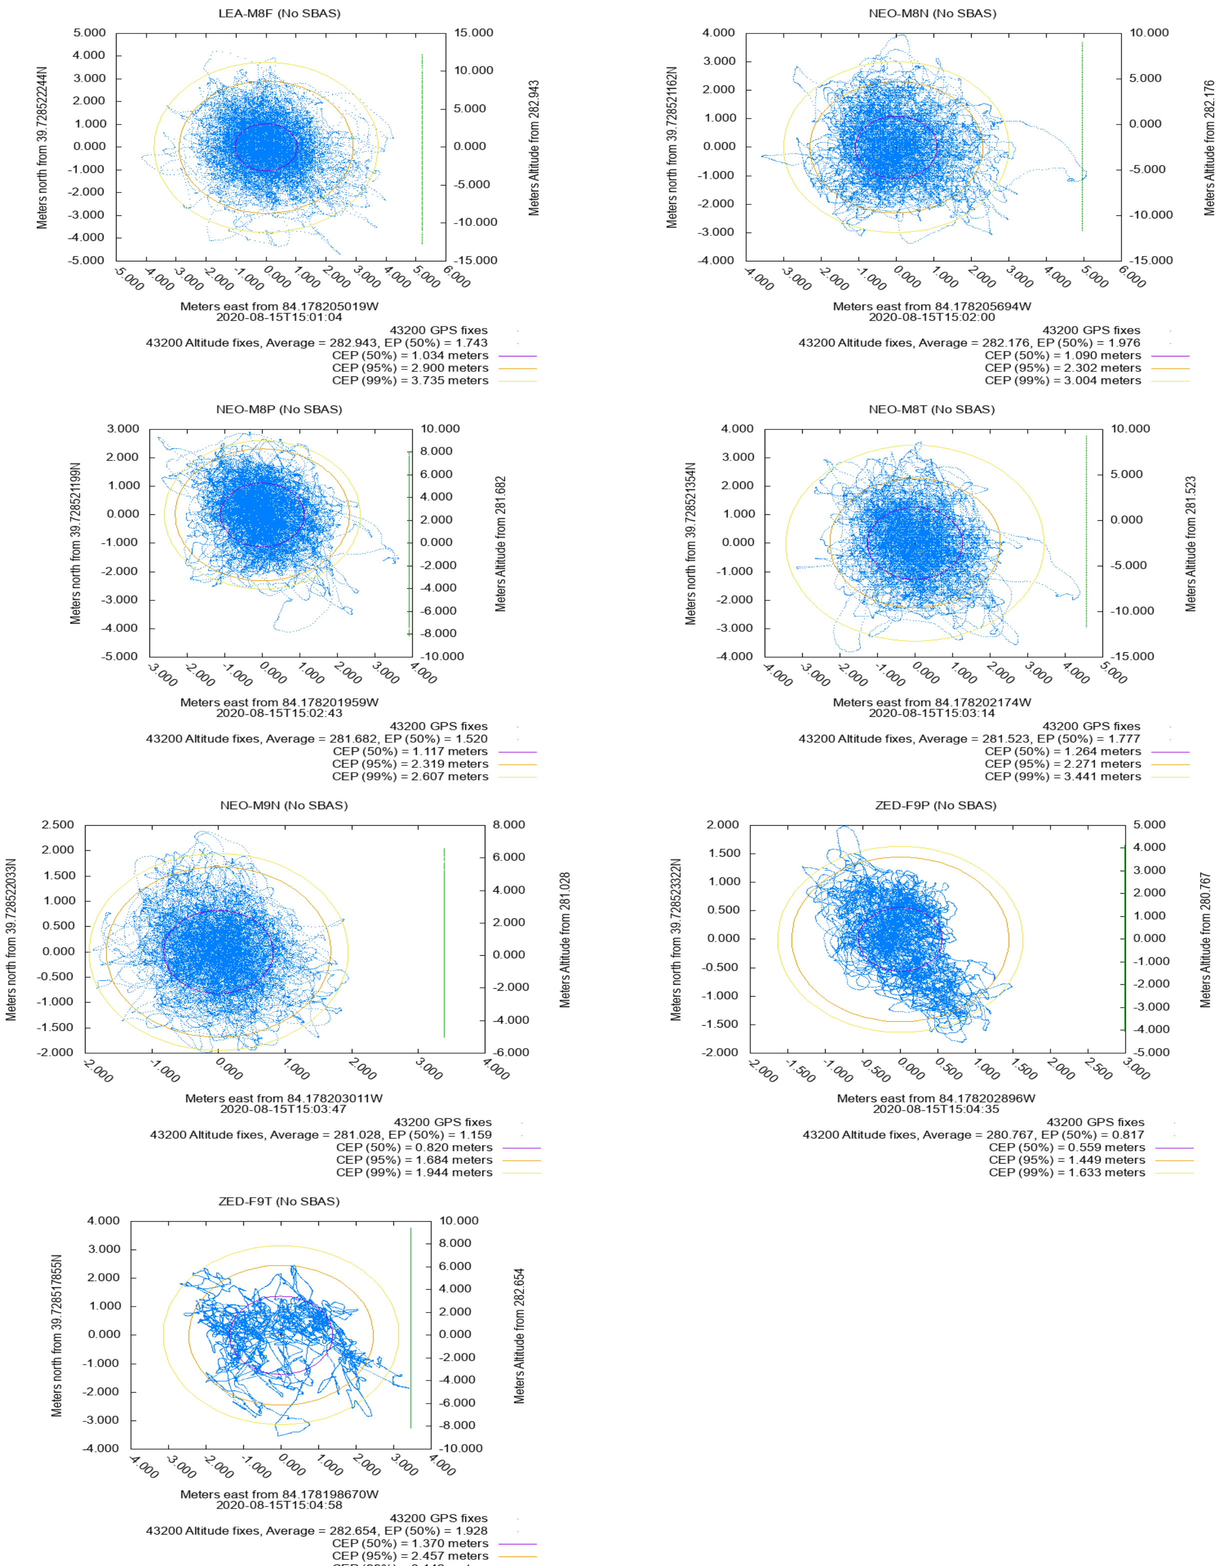

Figure 41: Circular Error Probability of Receiver Modules (via gpsprof)

Receiver	CEP – 50%	CEP – 95%	CEP – 99%	EP – 50%
LEA-M8F	1.034	2.900	3.735	1.743
NEO-M8N	1.090	2.302	3.004	1.976
NEO-M8P	1.117	2.319	2.607	1.520
NEO-M8T	1.264	2.271	3.441	1.777
NEO-M9N	0.820	1.684	1.944	1.159
ZED-F9P	0.559	1.449	1.633	0.817
ZED-F9T	1.370	2.457	3.142	1.928

Table 3: CEP/EP Values (via gpsprof)

For ease of comparison, Table 3 consolidates the CEP and EP results. It indicates that all the "8" series receivers perform similarly. The "9" series results are surprising. The single-frequency M9N performs significantly better than its M8N predecessor, and the ZED-F9P performs best of all the units, indicating that it uses a dual-frequency solution in its NMEA output. But the ZED-F9T, which one would think to have similar performance as the F9P, performs like the "8" series, and its scatter plot in Figure 42 looks quite different from that of the other receivers. Perhaps its timing firmware has not been updated to use a dual-frequency positioning solution, or there is some other issue.

Receiver	Latitude[32]	Longitude[33]	Altitude
LEA-M8F	39.xx8522244	-84.xx8205019	282.943
NEO-M8N	39.xx8521162	-84.xx8205694	282.176
NEO-M8P	39.xx8521199	-84.xx8201959	281.682
NEO-M8T	39.xx8521354	-84.xx8202174	281.523
NEO-M9N	39.xx8522033	-84.xx8203011	281.028
ZED-F9P	39.xx8523322	-84.xx8202896	280.767
ZED-F9T	39.xx8517855	-84.xx8198670	282.654
REFERENCE[34]	39.xx8519939	-84.xx8203511	281.521
AVERAGE	**39.xx8521309**	**-84.xx8202774**	**281.825**
MAX – MIN (m)	**0.607**	**0.602**	**2.176**
AVG – REF (m)	**0.152**	**-0.063**	**0.304**
MIN – REF (m)	**-0.231**	**-0.415**	**-0.754**
MAX – REF (m)	**0.376**	**0.187**	**1.133**

Table 4: Range of Average Positions

32 Values obfuscated to complicate ICBM targeting. At this latitude, 111029.3831 meters/degree.

33 Values obfuscated to complicate ICBM targeting. At this latitude, 85731.6808 meters/degree.

34 Reference position determined from 48 hour data collection using Trimble NetRS receiver, post-processed using IGS final data via NRCan (https://webapp.geod.nrcan.gc.ca/geod/tools-outils/ppp.php)

As shown in Table 4, the average positions calculated by gpsprof for the seven receivers shows a range in both latitude and longitude of about 0.6 meter, and a range of about 2.2 meters in altitude. The average of the seven receivers' average positions was 15 centimeters for latitude, six centimeters for longitude, and 30 centimeters for altitude away from the antenna location as determined by post-processing of results from a geodetic receiver.

8.2 RTK Positioning (NEO-M8P and ZED-F9P)

The "positioning" series receivers (NEO-M8P and ZED-F9P) have an internal Real Time Kinematic (RTK) processing engine. By providing an input stream of RTCM[35] messages from a base or reference station, the receivers will apply corrections that can result in centimeter or even better positioning accuracy.

As a test, the NEO-M8P and ZED-F9P were configured to use RTCM corrections fed via the u-blox "u-center" software[36] which has an NTRIP client capability.[37] Correction data was obtained from the Ohio Department of Transportation virtual reference station network[38] which has a reference station a few kilometers from the author's location. Data was collected from both receivers simultaneously. Due to time constraints and configuration issues it was only possible to collect data for the NEO-M8P with an RTK "FIX" solution for just under 8 hours. The ZED-F9P data run was more successful, but for comparability only data matching the time period collected from the NEO-M8P was analyzed.

As is shown in Table 4, the RTK processing engines work, providing significantly smaller CEP and EP than the autonomous measurements. The fact that the single- and dual-frequency receivers performed almost identically is surprising, but this may be a limitation of the reference network used. (The author currently has little experience configuring RTK systems.)

Receiver	CEP – 50%	CEP – 95%	CEP – 99%	EP – 50%
NEO-M8P	0.010	0.029	0.035	0.021
ZED-F9P	0.013	0.025	0.033	0.016

Table 4: Internal RTK Engine CEP/EP

The offsets in position shown in Table 5 are not what would be expected. Again, configuration issues may affect these results and further work needs to be done. Figure 43 is a scatter plot showing the spread of the data sets.

35 Real Time Correction Message protocol
36 https://www.u-blox.com/en/product/u-center
37 https://en.wikipedia.org/wiki/Networked_Transport_of_RTCM_via_Internet_Protocol
38 http://www.dot.state.oh.us/Divisions/Engineering/CaddMapping/Survey/VRS/Pages/default.aspx

Receiver	Latitude[39]	Longitude[40]	Altitude
NEO-M8P	39.xx8511818	-84.xx8194318	282.942
ZED-F9PN	39.xx8504462	-84.xx8179344	281.912
REFERENCE[41]	39.xx8519939	-84.xx8203511	281.927
RANGE (m)	**0.817**	**1.284**	**0.030**
ZED-F9P – REF (m)	**-1.718**	**-2.072**	**0.392**
NED-M8p – REF (m)	**-0.092**	**-0.788**	**0.421**

Table 5: Internal RTK Engine Range of Positions

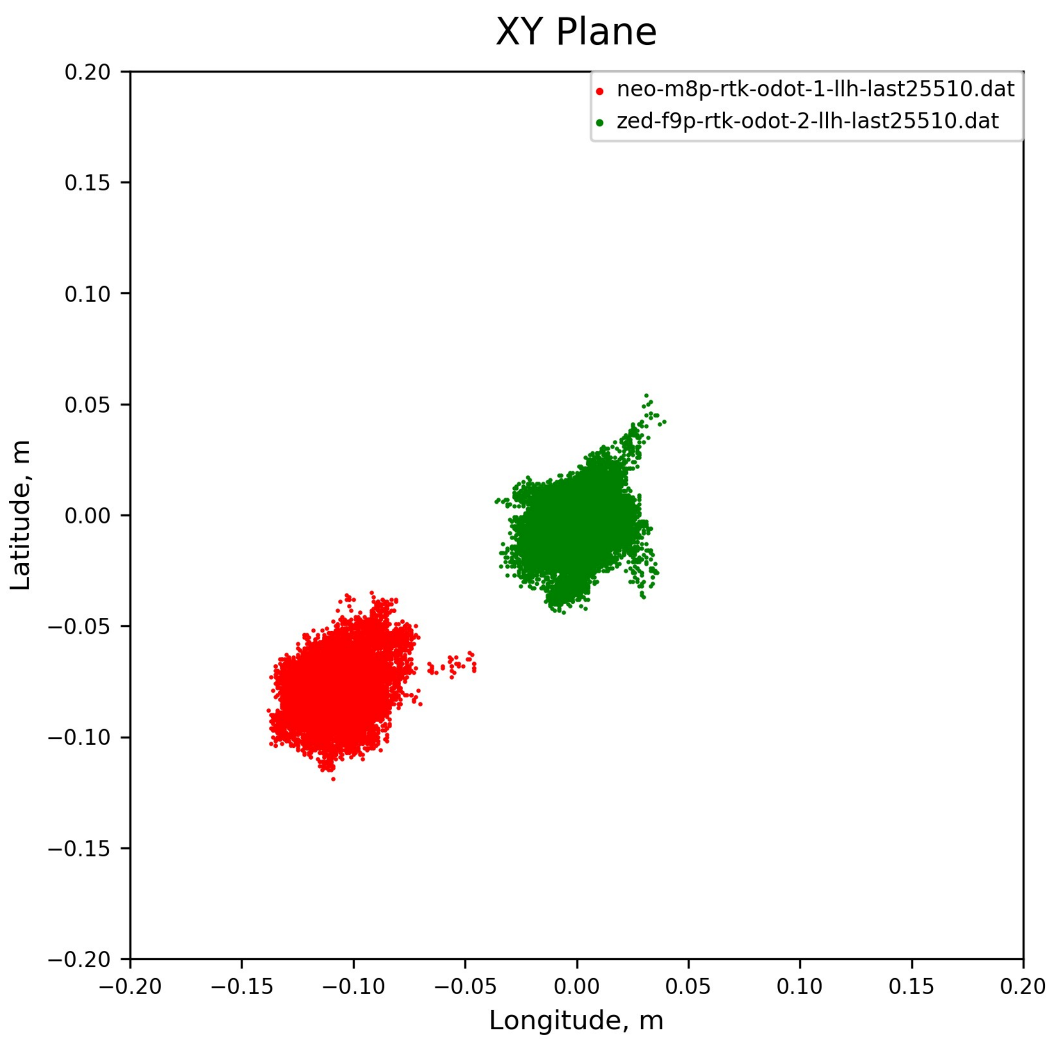

Figure 43: RTK Scatter Plot

39 Values obfuscated to complicate ICBM targeting. At this latitude, 111029.3831 meters/degree.
40 Values obfuscated to complicate ICBM targeting. At this latitude, 85731.6808 meters/degree.
41 Reference position determined from 48 hour data collection using Trimble NetRS receiver, post-processed using IGS final data via NRCan (https://webapp.geod.nrcan.gc.ca/geod/tools-outils/ppp.php)

8.3 Post-Processing

The greatest absolute positioning accuracy is generally achieved by post-processing raw satellite observation data using either a double-difference method employing corrections from known reference stations (such as is done by the Online Positioning User Service, or OPUS, web site operated by NOAA in the United States[42]) or the Precise Point Positioning method that applies correction information based on observed satellite orbital and clock information (such as is done by the service offered by Natural Resource Canada at https://webapp.geod.nrcan.gc.ca/geod/tools-outils/ppp.php?locale=en).

Raw data was recorded simultaneously from the NEO-M8P and ZED-F9P using GPS only, and from the ZED-F9T using both GPS and GLONASS. After allowing several weeks for the final IGS correction data to become available, the data from each receiver was submitted to the NRCan web site. For reference, data from a Trimble NetRS (dual frequency, GPS only) receiver was also submitted to NRCan for processing.

	NetRS (GPS)		NEO-M8P (GPS)		ZED-F9P (GPS)		ZED-F9T (GPS+GLONASS)	
	24 Hour	Sigma (95%)	24 Hour	Sigma (95%)	24 Hour	Sigma (95%)	24 Hour	Sigma (95%)
LAT ITRF2014	39 xx 42.67100	0.0068	39 xx 42.66852	0.3601	39 xx 42.67067	0.0090	39 xx 42.67086	0.0048
LON ITRF2014	-84 xx 41.53109	0.0124	-84 xx 41.53533	0.4131	-84 xx 41.53164	0.0160	-84 xx 41.53226	0.0084
EL HGT ITRF2014	247.101	0.0247	247.21	0.6522	247.1254	0.0370	247.1548	0.0217

Table 6: PPP Results from NRCan[43,44]

Table 5 shows the results. As can be seen, the ZED-F9P and ZED-F9T performed comparably to the NetRS. In fact, the ZED-F9T was slightly better. There is a reason for this.

As noted in Section 2.3, the ZED-F9 series receivers are dual-frequency, but they decode only the L2C signal that is present on GPS satellites launched since 2005. There are still some satellites in operation that do not support L2C, and the ZED-F9 receivers receive only the L1 signal from those. The much older NetRS receiver, on the other hand, receives the legacy L2P(Y) as well as the L2C signal.

As a result, when constrained to the GPS constellation, the ZED-F9 receivers have fewer available L2 satellites to use, perhaps increasing dilution of precision, and that slightly reduces their positioning accuracy. By adding the GLONASS constellation, the receiver has more usable satellites in view, and its performance is thereby somewhat improved, as is the case for this ZED-F9T test.[45] Since the NetRS receiver does not support GLONASS, it now loses out to the ZED-F9's ability to use satellites from both constellations. This discrepancy will diminish as more L2C-capable satellites are launched.

42 https://www.ngs.noaa.gov/OPUS/

43 Values obfuscated to complicate ICBM targeting.

44 These elevations are quite different from those shown in other tables and plots. Those are referenced to the NAD88 datum which more-or-less represents mean sea level, while ITRF2014 is based on the ellipsoid. As is shown, the difference can be tens of meters. GPS elevation measurements are not straight-forward.

45 Note that for timing as opposed to positioning applications, it is better to use a single satellite constellation in order to avoid multiple and slightly different clock systems.

9. CONCLUSION

The u-blox ZED-F9 series receivers represent a new generation of affordable dual-frequency GNSS hardware. The dual-frequency approach allows significant improvements in both timekeeping, with improved short and long term frequency stability, and positioning.

The quantization error or “sawtooth” correction capability of some u-blox receivers allows significant reduction in short term timing noise, and the 0D navigation mode improves long-term timing performance..

While the TIMEPULSE output of some receivers may be set to RF frequencies, the spectral purity of the output at those frequencies does not lend itself for use as a clock source in radio frequency systems.

The external interrupt (“EXT INT”) input of some receivers does not seem to be suitable for use as a precision timestamping counter due to unexplained glitches in the data output. The cause (and cure) of these glitches is as yet unknown.

The autonomous positioning performance of the receivers is generally similar, with the NEO-M9N and ZED-F9P having better CEP than the other units. The positioning receivers (NEO-M8P and ZED-F9P) can provide centimeter level results using their internal RTK engines and a nearby reference station. The dual-frequency receivers (ZED-F9P and ZED-F9T) can provide post-processed results similar to those of survey receivers.

APPENDIX 1
Receiver Models/Firmware Versions Tested

MODEL	FIRMWARE	PROT	HiPREC	qERR	SBAS	RAW	RTK	L2	PRICE
LEA-M8F	2.01	16.00	NO	NO	YES	NO	NO	NO	$99.00
NEO-M8N	SPG-3.01	18.00	NO	NO	YES	NO	NO	NO	$25.00
NEO-M8P-2	HPG 1.40REF	20.30	YES	YES	NO	YES	YES	NO	$149.00
NEO-M8T	TIM 1.10	22.00	NO	YES	YES	YES	NO	NO	$89.00
NE0-M9N	SPG 4.00	32.00	YES	NO	YES	NO	NO	NO	$25.00
ZED-F9P[46]	HPG 1.11	27.10	YES	YES	NO	YES	YES	YES	$199.00
ZED-F9T	TIM 2.01	29.00	YES	YES	YES	YES	NO	YES	$199.00[47]

FIRMWARE	=	Firmware version of unit tested
PROT	=	Protocol version
HiPrec	=	Seven-decimal-place NMEA output option
qERR	=	Availability of Quantization Error correction message
SBAS	=	SBAS reception/correction
RAW	=	Raw (pseudorange, doppler, carrier phase) output message
RTK	=	Internal RTK processing engine
L2	=	Dual Frequency L1/L2 receiver (L2C only)
PRICE	=	Quantity 1 price per u-blox.com visited 15 August 2020

46 HPG v. 1.13 firmware, which adds SBAS support and has other improvements, became available in June, 2020. The tests conducted for this paper began before this update was available.
47 In quantity volumes, the F9T unit becomes less expensive than the F9P.