The effects of the January 2016 UTC offset anomaly on GPS-controlled clocks monitored at NIST

Jian Yao, Michael A. Lombardi, Andrew N. Novick, Bijunath Patla, Jeff A. Sherman, and Victor Zhang
Time and Frequency Division, National Institute of Standards and Technology, Boulder, Colorado, USA

ABSTRACT

Errors in the Coordinated Universal Time (UTC) offset parameters broadcast by Global Positioning System (GPS) satellites caused many thousands of GPS-controlled clocks to be in error by approximately -13 µs on January 25-26, 2016. The erroneous UTC offset information was broadcast by 15 GPS satellites, or half of the available constellation, during the anomaly.

This paper discusses the technical reasons for the UTC offset anomaly, its effect on clocks controlled by GPS satellites (hereafter referred to as GPS clocks), and how it was detected at the National Institute of Standards and Technology (NIST). It then discusses the impact of the UTC offset anomaly on GPS clocks maintained and monitored by NIST, including clocks located in Boulder, Colorado and those at remote locations. Some analysis is presented of the effects of the UTC offset anomaly on GPS clocks located at 19 sites in North, Central, and South America.

I. INTRODUCTION

Through its remote frequency and time calibration services and its participation in the Sistema Interamericano de Metrologia Time Network (SIMTN) [1], NIST continuously monitors more than 90 single-frequency (L1 band) GPS receivers that are deployed as clocks, meaning that their function is to output a UTC time stamp and a 1 pulse per second (pps) time signal synchronized to UTC. In most cases, these clocks are part of a common-view system that is used to either measure or adjust the primary frequency or time standard at each location with respect to UTC(NIST), the official NIST time scale. Some of the GPS clocks are located at the NIST laboratories in Boulder, Colorado, but most, some 76 clocks, are located at various sites around the world, primarily in North, Central, and South America, as shown in Fig. 1.

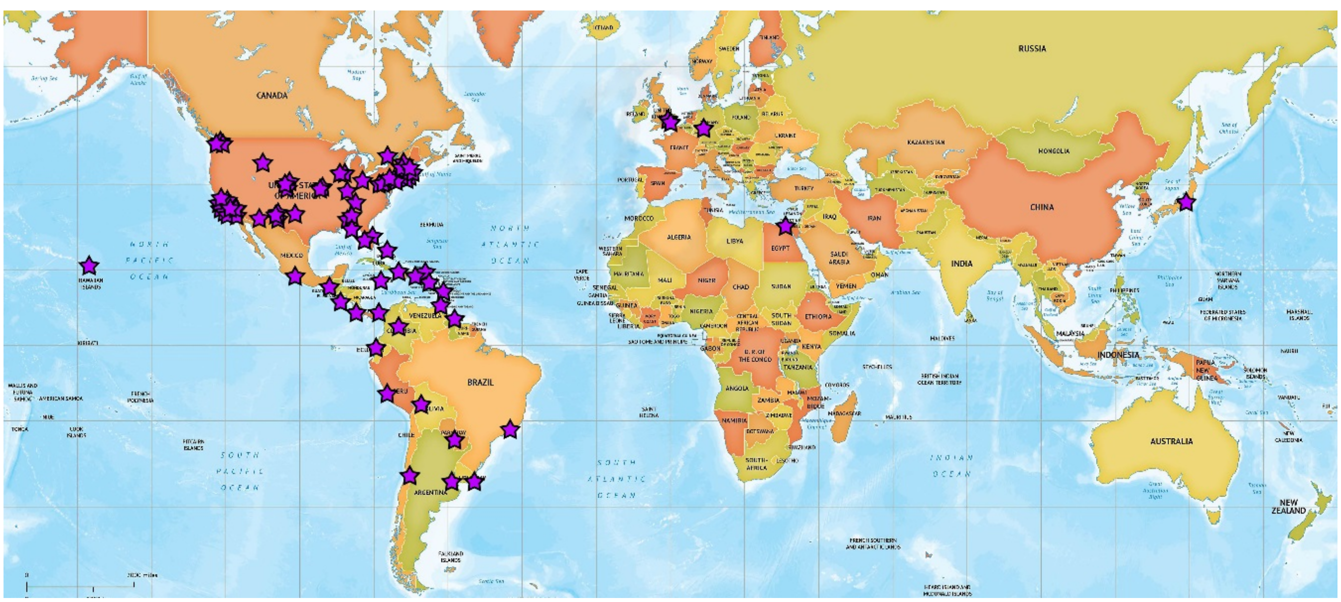

Figure 1. The locations of GPS clocks that are continuously monitored by NIST.

The large group of GPS clocks that NIST monitors and maintains allowed us to study the effects of the UTC offset anomaly that occurred on January 25-26, 2016, when, ironically, most of the authors of this paper were attending the 2016 Institute of Navigation Precise Time and Time Interval (ION PTTI) meeting. Here we present the results of our study. Section II provides a technical description of the UTC anomaly. Section III describes how the problem was detected by NIST personnel. Section IV summarizes the impact the UTC offset anomaly had on GPS clocks monitored and maintained by NIST both in Boulder, Colorado and at the sites shown in Fig. 1. Section V provides an analysis of the failure using data collected from 19 sites in North, Central, and South America. Finally, Section VI provides a summary.

II. TECHNICAL DESCRIPTION OF UTC OFFSET ANOMALY

The UTC offset anomaly was caused by a software "bug" that was triggered by the decommissioning of PRN 32 on January 25, 2016 at 2200 UTC [2], a satellite whose space vehicle (SV) number was 23. The PRN numbers refer to the satellite's pseudo random noise code, the SV numbers are assigned consecutively in the order of the satellite launch dates. The software failure occurred not because the highest PRN number (32) was removed from the constellation, but rather because the lowest SV number (and hence the oldest satellite) was removed. The first of the Block IIA satellites, SVN 23 was launched on November 26, 1990, and thus was in service for more than 25 years.

The software failure caused three of the parameters used to correlate Coordinated Universal Time (UTC) to GPS time, A_0, t_{ot}, and WN_t, to be incorrect [3]. These parameters are included in subframe 4, page 18, of the GPS navigation message. The content of subframes 4 and 5 is common for all satellites so a receiver can obtain the UTC parameters by tracking just one satellite. The UTC offset correction, Δt_{UTC}, is computed as [4]

$$\Delta t_{\mathrm{UTC}} = \Delta t_{\mathrm{LS}} + A_0 + A_1\,(t_{\mathrm{E}} - t_{\mathrm{ot}} + 604800(WN - WN_{\mathrm{t}})), \qquad (1)$$

where

Δt_{LS} is the number of leap seconds introduced into UTC since GPS time began (equal to 17 when the anomaly occurred),

A_0 is the constant UTC offset parameter expressed in seconds,

A_1 is a dimensionless frequency offset value that allows the correction of the time error accumulated since the UTC reference time, t_{ot}, which is when A_0 was last determined,

t_E is GPS time (also known as the time of interest or the time being converted to UTC),

604800 is a constant that equals the number of seconds in one week.

t_{ot} is the reference time for UTC data,

WN is the GPS week number, and

WN_t is the UTC reference week number.

The first part of Eq. (1), $\Delta t_{LS} + A_0$, essentially takes care of the UTC correction. The Δt_{LS} term is the large, integer second part of the correction, equal to the number of leap seconds that have occurred since January 6, 1980, the start of the GPS time scale. The Δt_{LS} value was correctly broadcast as 17 during the anomaly. The A_0 term is the small, nanosecond part of the correction, equal to the difference between the GPS and UTC(USNO) second markers, with UTC(USNO) being the UTC time scale maintained by the United States Naval Observatory. The A_0 term is broadcast in units of seconds, but is typically $< 1 \times 10^{-8}$ s, or < 10 ns. For example, the three values of A_0 broadcast by GPS immediately before and after the UTC anomaly (six values total) ranged from -2.79 ns to +1.86 ns. The erroneous value of A_0 broadcast during the timing anomaly was approximately -13696.03 ns, which rounds to the -13.7 µs value [5] that was widely reported by media outlets.

The second part of Eq. (1) fine tunes the UTC output of a GPS clock by applying a dimensionless frequency offset, provided by A_1, as a drift correction for the interval between the time specified by t_{ot} and WN_t and the current time. This is normally a sub-nanosecond correction, because A_0 is updated in the GPS broadcast more than once per day and the drift correction supplied by A_1 is typically near 1 ns per day. During the anomaly, a correct A_1 value of 1.24345×10^{-14} was broadcast, a frequency offset that translates to a time offset of about 1.07 ns per day. However, this correction was applied across a much longer than normal interval. The correct GPS week number was 1881 and the correct value for WN_t, broadcast in an 8-bit field, should have been modulo (1881, 256) or 89. Instead, it was broadcast as 0. Similarly, t_{ot} is also broadcast as an 8-bit field. The correct value would have expressed a time, given in multiples of 4096 s (2^{12} s), that occurred approximately 70 h after the first valid transmission time for the UTC data set [4]. Instead, t_{ot} was also erroneously broadcast as 0 during the anomaly.

Referring to Eq. (1), broadcasting WN_t as 0 resulted in $WN - WN_t$ being equal to 623 days (89 weeks × 7 days). Broadcasting t_{ot} as 0 resulted in $t_E - t_{ot}$ being equal to ~2 days, instead of the usual fractional part of a day, because the anomaly occurred near the Monday/Tuesday transition, about two days after the GPS week began on the Saturday/Sunday transition. Therefore, the drift correction supplied by A_1 was applied to a period of ~625 days (623 + ~2 days) [3, 5]. This added a correction of ~671 ns and reduced the UTC offset to about -13024 ns. Thus, GPS clocks that applied A_1 were offset from UTC by about -13.0 µs and those that did not were offset by about -13.7 µs [5]. The GPS clocks in our study all applied the A_1 drift correction.

The erroneous UTC offset parameters were transmitted by 15 of the 30 satellites that remained in the GPS constellation after PRN 32 was decommissioned. Table 1 lists the 15 satellites and the periods when bad data were transmitted [5].

Table 1. Periods when erroneous UTC correction data were broadcast by GPS satellites.

PRN	Start of bad transmission	End of bad transmission	Duration (hh:mm)
13	1/25/16, 23:30	1/26/16, 11:45	12:15
9	1/25/16, 23:45	1/26/16, 12:45	13:00
10	1/25/16, 23:45	1/26/16, 13:15	13:30
29	1/26/16, 00:15	1/26/16, 12:30	12:15
14	1/26/16, 00:30	1/26/16, 11:45	11:15
27	1/26/16, 00:45	1/26/16, 12:45	12:00
6	1/26/16, 01:00	1/26/16, 12:30	11:30
11	1/26/16, 01:15	1/26/16, 12:15	11:00
20	1/26/16, 03:00	1/26/16, 12:00	09:00
3	1/26/16, 03:15	1/26/16, 13:15	10:00
25	1/26/16, 04:30	1/26/16, 12:15	07:45
28	1/26/16, 05:30	1/26/16, 12:00	06:30
24	1/26/16, 07:00	1/26/16, 13:15	06:15
26	1/26/16, 07:00	1/26/16, 13:00	06:00
23	1/26/16, 08:00	1/26/16, 13:00	05:00

The time stamps in Table 1 refer to GPS system time. The table indicates that the periods when bad data was broadcast by the satellites ranged from 5 h for PRN 23 to 13 h, 30 min for PRN 10. Nine of the satellites were in error for at least 10 h, the remaining six were in error for 9 h or less. The duration of the UTC anomaly was 13 h and 45 min, beginning when PRN 13 first sent bad data on January 25th at 23:30 and ending when PRNs 3 and 24 stopped sending bad data on January 26th at 13:15.

The interval between the last bad upload to the satellites (07:46:18, to PRN 23) and the first corrected upload (11:43:48, to PRNs 13 and 14) was 3 h, 57 min, and 30 s. After the first corrected upload the remaining 13 satellites were corrected within 1 h, 27 min, and 30 s [5].

III. DETECTION OF UTC OFFSET ANOMALY BY NIST PERSONNEL

NIST personnel were first alerted about the UTC offset anomaly by an alarm system that was not intended to monitor GPS, but instead was designed to monitor the UTC(NIST) time scale in Boulder, Colorado. This system has been continuously operational since 2010 and works by examining near real-time measurement data collected by the SIMTN. The alarm system continuously compares UTC(NIST) via common-view GPS measurements to seven independent time scales, with new comparison results available every 10 minutes. On the day of the anomaly, the seven time scales serving as verification or check standards were the national time scales of Argentina, Brazil, Canada, Mexico, and Panama, the secondary UTC(NIST) time scale located in Fort Collins, Colorado, and a hydrogen maser time scale maintained by a NIST customer in Massachusetts. The alarm system is based on a 50 ns threshold, but it is not uncommon for one or two of the seven check standards to differ from UTC(NIST) by more than 50 ns at any particular time. For this reason, an alarm is only triggered if four of the seven time scales show a difference of more than 50 ns, a rare condition that indicates that it is highly likely that UTC(NIST) is out of tolerance. When an alarm occurs, text messages are immediately sent to NIST staff members and continue to be sent at 10 minute intervals until the time scale is restored to normal operation.

On January 25, 2016, the NIST staff members who are connected to the alarm system were in Monterey, California for the 2016 ION PTTI meeting. They received the first text message from the alarm system at 5:10 p.m. Pacific Standard Time (PST), or 01:10 UTC, while the PTTI meeting was still in progress. By then, as indicated in Table 1, seven GPS satellites were transmitting bad UTC data and three had been doing so for more than an hour.

The alarm system did not trigger earlier for two reasons. The first reason is that a satellite transmitting bad data would not only have to be received by a common-view GPS receiver that was part of the alarm network, but it would also have to be the particular satellite that was supplying the UTC offset information to that receiver in order to cause a GPS clock failure. As noted previously, the same UTC offset information is available from all of the satellites, therefore a receiver can obtain it by decoding the broadcast from just one satellite. The receivers in our study refresh this information every 12.5 min (750 s), which is the period required to receive a new GPS navigation message, even though new UTC information is only uploaded to the satellites about four times per day. The uploads to the satellites are sent sequentially, so some satellites will temporarily have newer information than others. In most cases, the receiver will obtain the UTC parameters from the same satellite that it utilized during the previous refresh, and most of the time these will be the same parameters that it already had. The GPS clock's parameters will only change if the satellite it has been receiving has just been updated, or if it switches to a different satellite that has newer information. The exact criteria the clocks in this study used to switch satellites is unknown, but it appears that the elevation angle of the dropped satellite was always decreasing, whereas the elevation angle of the newly selected satellite was always increasing. The second reason for the delayed reporting of the problem was that the SIMTN, the alarm system, and all of the NIST remote calibration services that distribute UTC(NIST), are based on the common-view technique. This means that if bad data was sent to both of the receivers participating in a particular common-view link, then the link would still be good. The link would only produce bad data if just one, and not both, of the receivers were bad.

Table 2. The GPS satellites broadcasting erroneous UTC parameters when the first alarm message was sent to NIST staff.

PRN	Started broadcast of erroneous UTC offset information	Was PRN being received by NIST alarm system receiver	
		When bad data were first uploaded to satellite?	When first alarm was sent?
13	1/25/16, 23:26:24	No	No
9	1/25/16, 23:38:54	Yes, from 22:30	Yes
10	1/25/16, 23:38:54	No	No
29	1/26/16, 00:03:54	No	No
14	1/26/16, 00:16:24	No	No
27	1/26/16, 00:41:24	Yes, from 23:40	Yes
6	1/26/16, 00:53:54	No	No
11	1/26/16, 01:06:24	No	No

Our systems did not store the UTC offset information, but Table 2 indicates that the erroneous data that caused the initial alarm message was sent by either PRN 9 or PRN 27. The NIST alarm system was receiving seven satellites (PRNs 7, 8, 9, 16, 23, 27, and 30) when the first alarm was sent and only PRNs 9 and 27 were in error. When the alarm occurred the receiver's time offset with respect to UTC(NIST) moved from near -10 ns to ~13011 ns as shown in Table 3. This approximate -13021 ns time step agrees with the -13024 ns noted earlier when A_1 was applied, with the few nanoseconds of difference being within the uncertainty of the measurement.

Table 3. The UTC via GPS – UTC(NIST) measurements recorded before and after the first alarm message.

Date	Time (UTC)	UTC via GPS – UTC(NIST), ns
1/26/16	00:50	-10.00
1/26/16	01:00	-10.24
1/26/16	01:10	4737.98*
1/26/16	01:20	13011.14
1/26/16	01:30	13010.28

* the receiver refreshed the UTC offset parameter during this 10-minute segment so only the last part of this segment was bad. However, the average value for the segment exceeded the 50 ns threshold and triggered the alarm messages.

The NIST alarm receiver showed a ~13 µs offset for at least parts of 43 of the 83 10-minute segments recorded during the UTC offset anomaly. Alarm text messages were actually sent during 48 of the 83 segments. The segments when text messages were sent did not necessarily correspond to the segments when the NIST alarm receiver was in error, hence the difference between 43 NIST receiver errors and 48 alarm conditions. There were multiple segments when the error was not seen at NIST but was seen by at least four of the other seven GPS receivers involved in the common-view comparisons. There were also segments when NIST and at least four of the other seven receivers showed the same ~13 µs error (an error that was cancelled by the common-view technique), so the alarm did not trigger.

As previously noted, the alarm system was designed to report UTC(NIST) errors, and not GPS errors. Therefore, when the alarm messages were first received, we initially thought that UTC(NIST) was the source of the failure. However, after a large amount of measurement data from our network of GPS clocks was analyzed from Monterey via the Internet (the ~13 µs offset was found at many sites with some showing it as early as 23:40 UTC on January 25th), and after a number of tests of the UTC(NIST) time scale were performed by one of the authors who remained at NIST in Boulder, it became obvious that GPS was the source of the failure. We then reported the error by telephone to the GPS operations center at Schriever Air Force Base in Colorado between 8 and 9 p.m. Pacific Standard Time (between 04:00 and 5:00 UTC on January 26th). A written report was sent to the United States Coast Guard Navigation Center (NAVCEN) through their web portal and by email at 11:02 p.m. PST (07:02 UTC). The email indicated that the offset was "almost exactly" 13 µs and that it was intermittently occurring at many of the sites shown in Fig. 1. NIST is believed to have been the first organization to report the error to the GPS operations center.
Alarm text messages continued to be sent intermittently throughout the night. During the night, we were, of course, not sure what was causing the GPS problem or if or when it would stop. We received the last alarm message at 4:20 a.m. Pacific Standard Time (12:20 UTC). Less than one hour later we received a phone call from NAVCEN explaining that the problem had been corrected. Official notification from the United States Air Force was published via email and the Internet on the evening of January 27th [6].

IV. IMPACT OF THE UTC OFFSET ANOMALY ON GPS CLOCKS MONITORED BY NIST

On the day of the UTC offset anomaly, there were 92 GPS clocks being maintained and/or monitored by NIST that were based on 8-channel or 12-channel single frequency (L1 band) receivers. Figure 1 shows the location of 76 of these clocks, the remainder were located at the NIST laboratories in Boulder, Colorado or at temporary field sites. For various reasons, including systems that were down for maintenance, and data files that could not be recovered due to hardware, software, or Internet failures at field sites, we could not collect data from 12 of the 92 clocks, leaving us with 80 different GPS clocks whose measurements we could examine. Seven of these clocks, or 8.8 %, ran continuously throughout the UTC offset anomaly (recording measurement data every second) without recording any ~13 µs time steps. Two more GPS clocks did not have any time steps, but did have periods during the anomaly when they stopped recording measurements (one stopped recording for 30 minutes, the other for 50 minutes). We believe, but cannot say conclusively, that those data outages were caused by issues other than the UTC offset anomaly. If those two clocks are counted as unaffected, then 11.3 % of the GPS clocks made it through the event unscathed.

It may seem improbable that any GPS clock that applied the UTC correction would be able to make it through an event of this duration without being affected. For example, consider that the UTC anomaly lasted for 13 h, 45 m, and that each clock refreshed its almanac every 12.5 m, or 66 times during the anomaly. If we were to assume that each GPS clock had a 50 % probability of receiving a bad satellite during each of its 66 attempts to get a new almanac then the odds of never receiving a bad satellite would be the same as the odds of flipping a coin 66 times and never getting tails, or a staggering 7.4×10^{19} to 1 (2^{66} to 1). Even if only one of the eight satellites being received by an 8-channel device were bad during each of the 66 refreshes; in other words, if the probability of getting a good satellite was 7/8 or 87.5 %, then the probably of making it through the anomaly without utilizing a bad satellite would be just $(7/8)^{66}$, or 0.00015 %.

In reality, however, when the GPS clocks refresh their UTC parameters they usually obtain the information from the same satellite that was used for the last refresh, and that particular satellite's almanac had probably not changed because it is only updated about four times per day. This means that the same satellite usually provides the same UTC offset data to a GPS clock for multiple hours. Thus, as it turns out, most of the GPS clocks "flipped a coin" not 66 times, but only three or four times during the UTC anomaly. The probability of either heads or tails appearing four consecutive times would be $(1/2)^4$ or 6.25 %, and the probability of three consecutive flips of either heads or tails would be $(1/2)^3$ or 12.5 %. This analysis is probably far too simple, but the range from 6.25 % and 12.5 % overlaps the range of what actually occurred. For completeness, Table 4 shows the locations of the nine GPS clocks that defied the odds and did not record any time steps during the UTC anomaly.

Table 4. The locations of nine GPS clocks monitored by NIST that were not affected by the UTC anomaly.

Location	Latitude	Longitude
Querétaro, Mexico	20° 32’ 13.4" N	100° 15’ 17.2" W
Ottawa, Canada	45° 27’ 14.9" N	75° 37’ 25.8" W
Rio de Janeiro, Brazil	22° 53’ 44.6" S	43° 13’ 27.4" W
Cairo, Egypt	30° 6’ 59.0" N	30° 54’ 55.0" E
Waterton, Colorado	39° 29’ 50.9" N	105° 5’ 52.2" W
El Segundo, California	33° 54’ 51.0" N	118° 23’ 21.3" W
Aurora, Illinois	41° 47’ 45.9" N	88° 14’ 37.0" W
Panama City, Panama*	9° 0’ 12.3" N	79° 35’ 3.6" W
San Jose, Costa Rica#	9° 55’ 52.9" N	84° 3’ 22.6" W

* failed to record 50 minutes of measurement data
# failed to record 30 minutes of measurement data

From the NIST perspective there were no critical problems caused by the UTC offset anomaly, but there were numerous service disruptions; and the situation had to be explained to each customer that it affected. Nearly 50 of the GPS clock locations shown in Fig. 1 were the sites of NIST remote calibration customers who receive monthly calibration reports. These reports had to be annotated to explain what had happened, and a few calibrations that were in progress at customer sites had to be restarted.

Some NIST customers have NIST disciplined clocks and oscillators that lock to UTC(NIST) using near real-time common-view data. These devices lost lock, and in a few cases, took nearly a day to relock. This affected NIST’s financial market customers as described in Ref. [7], although their clocks still remained well within their acceptable tolerances. A few GPS disciplined oscillators operated at the NIST laboratories in Boulder also lost lock, with at least one taking several days to relock.

Nearly all of GPS clocks discussed in this study were deployed in common-view systems, so it is fair to ask why the UTC offset parameters were being applied. It is true that utilizing the UTC offset parameters is not required for common-view GPS measurements (the 1 pps signals do not need to be synchronized with UTC), and thus many common-view systems, including the primary common-view receivers operated by NIST, do not apply the A_0 and A_1 corrections and were not affected by the anomaly. The GPS clocks involved in this study, however, were part of systems that time-tagged measurements with UTC time-of-day information (HH:MM:SS). These systems read a time interval counter once per second, and therefore a UTC time tag that only included the leap second correction, Δt_{LS}, and that excluded A_0 and A_1 would be sufficiently accurate for time tag purposes and immune to the UTC anomaly. However, these GPS clocks do not allow partial use of the UTC offset parameters, thus it was not possible (with the standard firmware command set) to apply the leap second correction without also applying A_0 and A_1. When the systems were originally designed we elected to use all of the UTC offset parameters so we would not have to obtain leap second information from another source in order to generate UTC time tags.

V. ANALYSIS OF THE UTC OFFSET ANOMALY USING A SUBSET OF GPS CLOCKS

This section presents some analysis of 19 of the GPS clocks that were being monitored by NIST at the time of the UTC anomaly. The clocks in this subset were chosen because they were each being compared to independent frequency and time standards during the UTC anomaly. The independent standards were either in the form of ensemble time scales, hydrogen masers, or cesium clocks. Data from other sites was excluded for a variety of reasons. In some cases, the GPS clock was being utilized to discipline the reference clock; either directly to GPS or to UTC(NIST) through common-view, so the local time standard was not independent. In other cases, the local time standard was not particularly stable or accurate and therefore was excluded; for example, the reference clock at those sites might have been a rubidium clock. Finally, some data files were incomplete, unavailable, or in many cases, only contained data in the form of 1-hour averages, instead of the more useful 1-minute or 10-minute average formats. The locations of the 19 clocks chosen for the analysis are shown in Fig. 2.

Figure 2. The locations of 19 GPS clocks that were compared to independent time standards during the UTC anomaly.

Referring back to Table 1, we can obtain the periods when the various GPS satellites were broadcasting bad information. During the period from 08:00 to 11:45 on January 26, 2016, 50 % of the GPS satellites were bad. Ideally, we would expect that the probability of timing errors detected by the GPS clocks to reflect the shape of the curve shown in Fig. 3. However, Fig. 4 shows the actual percentage of errors detected from measurements of the 19 GPS clocks.

Figures 3 and 4 have almost the same start time and end time for the UTC error (i.e., 23:30 on January 25th, and 13:15 on January 26th). However, the probability of a time error is about 35 % during the period from 01:00 to 06:00 on January 26th, and lower, about 20 %, during the period from 08:00 to 11:45. This differs from the shape of Fig. 3, which indicates that 50 % of the satellites were bad from 08:00 to 11:45. These discrepancies may arise for three reasons. First, only 19 GPS clocks are used in the Fig. 4 data, which makes the statistical uncertainty large. Second, the locations of the GPS clocks are not distributed uniformly around the globe but are all located in the Americas. Third, as we mentioned in Section III, the GPS clocks do not randomly select a satellite every 12.5 min. Instead, they usually continue to obtain the UTC corrections from the same satellite until this satellite does not meet some criteria (for example, the elevation angle becomes too low or the signal becomes too weak). This produces effects similar to those of a low-pass filter. These factors help explain why the Fig. 4 data is not highly correlated with Fig. 3.

We analyzed the average autocorrelation of the time output by the 19 GPS clocks during the period from 08:00 to 11:45 on January 26th when 50 % of the satellites were bad. Figure 5 shows the autocorrelation for the first three hours of this period. From this analysis, we find that the time output of the clocks is still highly correlated even after three hours. If a receiver selected a random satellite every 12.5 min, the autocorrelation would be near 0 after a time delay of more than 12.5 min. The high autocorrelation in Fig. 5 further confirms our earlier statement that the GPS clocks continued to utilize the same satellite for the UTC correction, instead of randomly selecting a new satellite every 12.5 min.

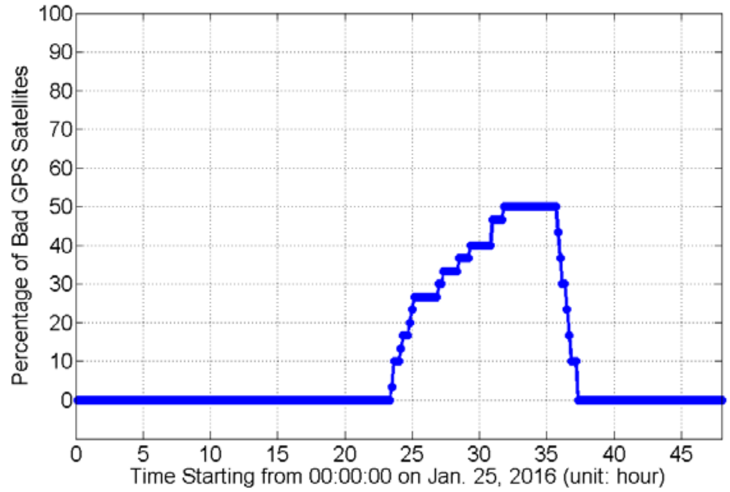

Figure 3. Percentage of bad GPS satellites vs. time.

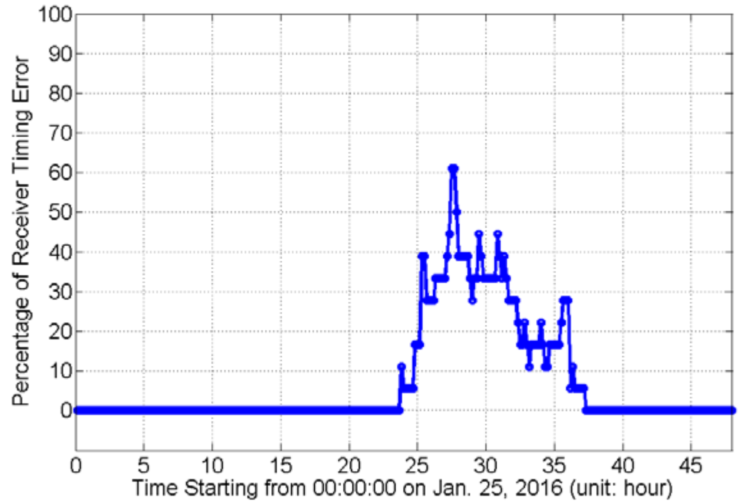

Figure 4. Percentage of GPS clock errors vs. time.

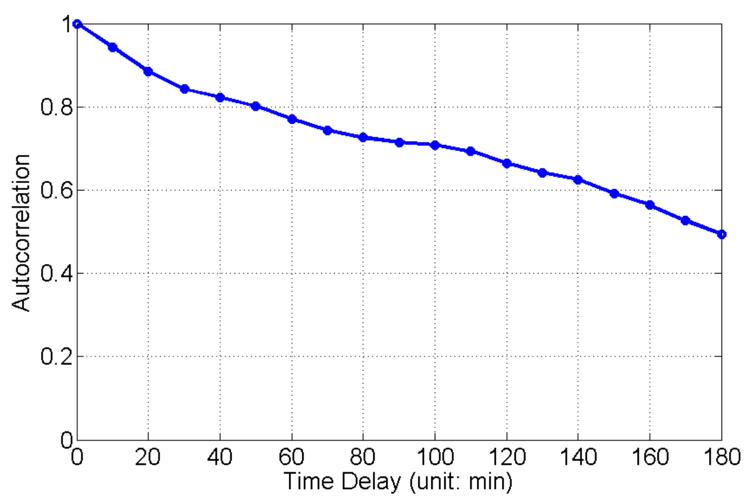

Figure 5. Autocorrelation vs. time delay. Data from 19 GPS clocks are used for the autocorrelation computation.

Using the Fig. 5 data we can derive the theoretical probability of a GPS clock not being affected during the entire UTC anomaly. The data show that the autocorrelation is 0.9426 at a time delay of 10 min. From the definition of autocorrelation of time series $\{X_1,\ X_2, \ldots, X_n\}$ then

$$R(k) = \frac{1}{(n-k)\sigma^2} \sum_{t=1}^{n-k} (X_t - \mu)(X_{t+k} - \mu), \tag{2}$$

where the true mean μ = (−13024.57 × 50 % + 0 × 50 %) ns = −6512.285 ns, and the deviation σ = 6512.285 ns, we can determine that the probability of X_{t+1} = 0 is 0.9713 when X_t = 0. In other words, P(X_{t+1}=0|X_t = 0) = 0.9713 and P(X_{t+1} = −13024.57|X_t = 0) = 0.0287. Of course, this probability only applies to the scenario when 50 % of the satellites are bad. The probability P(X_{t+1} = −13024.57|X_t = 0) should be proportional to the percentage of bad satellites. For example, if there is only one bad satellite, then P(X_{t+1} = −13024.57|X_t = 0) = 1/15 × 0.0287 = 0.0019 and P(X_{t+1} = 0|X_t = 0) = 0.9981. To assure that a receiver is not affected during the whole incident, we require all X_t to be 0, given the initial X_1 = 0. By multiplying the probability sequence, we can obtain the probability that a receiver would not be affected during the whole anomaly, or P(no-error) = 18.06 %. This is a theoretical estimation, based on the assumptions that the receivers under study have a uniform geographic distribution and that the autocorrelation coefficient in Fig. 5 is accurate. As mentioned in Section IV, 11.3 % of the GPS clocks monitored by NIST were not affected. To reiterate, the discrepancy between the actual and theoretical percentages is probably due to the fact that the full set of 80 GPS clocks were not uniformly distributed around the world (they were mainly located in the Americas) and that the subset of 19 GPS clocks used in the autocorrelation analysis is too small of a sample to eliminate statistical bias.

VI. SUMMARY

This paper has described the technical reasons for the January 2016 UTC offset anomaly and its effect on GPS clocks monitored by NIST. We also discussed how the NIST alarm system was able to detect this timing anomaly. In addition, we have analyzed the probability that a GPS clock would be affected by an anomaly of this nature. The problem could have been avoided by the use of GPS clocks that did not apply the UTC offset parameters, or that checked the GPS navigation message for errors; either by cross comparing data from multiple satellites or by discarding values that appeared to be outside of a reasonable range.

VII. ACKNOWLEDGEMENTS

The authors thank Edward Powers of the USNO and Stephen Hamilton of NAVCEN for supplying information and for patiently answering our questions. We also thank Stefania Romisch and Joshua Savory of NIST for helpful comments and corrections.

VIII. REFERENCES

[1] M. Lombardi, et al., 2011, "The SIM Time Network," *Journal of Research of the National Institute of Standards and Technology*, 116(2), 557–572.

[2] GPS Notice Advisory to Navstar Users (NANU), number 2016008, January 25, 2016.

[3] EndRun Technologies, *Field Service Bulletin 160126*, January 28, 2016.

[4] Global Positioning Systems Directorate, "Navstar GPS Space Segment/Navigation User Interfaces," *Interface Specification IS-GPS-200H*, 2013.

[5] K. Kovach, P. J. Mendicki, E. D. Powers, and B. Renfro, 2016, "GPS Receiver Impact from the UTC Offset (UTCO) Anomaly of 25-26 January 2016," *Proceedings of 29th ION GNSS+*, 9 p., Portland, Oregon, USA, September 2016.

[6] United States Air Force, "Air Force Official Press Release - GPS Ground System Anomaly," January 27, 2016.

[7] M. Lombardi, A. Novick, B. Cooke, and G. Neville-Neil, 2016, "Accurate, Traceable, and Verifiable Time Synchronization for World Financial Markets," *Journal of Research of the National Institute of Standards and Technology*, 121, 436-463.