GPS IIR RUBIDIUM CLOCKS: IN-ORBIT PERFORMANCE ASPECTS

Dr. Marvin Epstein
marvin.epstein@itt.com
+1(973) 284-2614

Gerald Freed
gerald.freed@itt.com
+1 (973) 284-2104

Dr. John Rajan
john.rajan@itt.com
+1 (973) 284-4835

ITT Industries Aerospace/Communications Division
NJ Engineering Center
100 Kingsland Rd.
Clifton, NJ 07014-1993, USA

Abstract

The GPS Block IIR rubidium clocks have proved to be the best performers in the current GPS constellation, starting with the first IIR operational clock (launched on July 23, 1997). This discussion covers a number of topics, including the frequency accuracy, the frequency variation, and the projected lifetimes of these clocks. The performance of these clocks in orbit and their life expectancy are compared with that experienced with previous GPS clocks. The superior performance and lifetimes of the GPS IIR Rb clocks are highlighted. Because these clocks are so well behaved, it has been possible to detect various features and phenomena that were not previously detectable. The full drift of the IIR clocks is not observed by the user because of onboard cancellation by the Time Keeping System. The diagnostic and re-programmability capability of the GPS IIR Satellites allows these discrepancies to be addressed. The long-lived features of the clocks indicate that they will outlive the GPS satellite mission duration.

INTRODUCTION

This paper presents an overview of the behavior of the Rubidium Atomic Frequency Standard (RAFS) on the Block IIR satellites. These clocks, as a family, have proved to be the best performers in the current Global Positioning System (GPS) constellation **[1]**. This is an interesting product, whose development started in 1980. ITT let a development contract in 1990 and a production contract in 1992 for the IIR clocks. They were tested and delivered between 1994 and 1997. Each IIR satellite contains three Rubidium Atomic Frequency Standard (RAFS) clocks. As of 2003-09-01, there are eight vehicles currently in orbit and 12 more vehicles ready to be launched as needed. The navigation mission of the GPS constellation is to provide a global source of navigation and time information. As a free automatic bonus, the IIR clocks provide an interesting experiment in the atomic clock design with 60 clocks, all of the same type and manufacture. These will be observed as a life test over decades. There is no problem with having these clocks monitored, because they will be observed by worldwide users due to their involvement with GPS. One can view this as a pure experiment with one design. This is in contrast to the earlier blocks of GPS satellites, which had a mixture of cesium- and rubidium-based clocks of various designs. The design of the RAFS was a new design for space deployment. Although initially there were concerns about this design, the results have been so successful that these clocks have become a benchmark for GPS clocks. Six of the eight operational clocks have significant history with between 2 and 5+ years of in-orbit operation. Also, two RAFS of the same design have been on life test at Naval Research Laboratory (NRL) for over 6 years. Although there are many more years of life ahead for the 60 IIR clocks, this is probably a good time to: 1) Review what is known about these clocks. 2) Examine the issues that have been raised by these observations. 3) Determine what might be done to improve the clock performance.

RESULTS

The results presented in this paper are a brief survey of the available information. More detailed information is available in the material referenced at the end of this paper. The bulk of the information in this paper is a collection of the work done by a wide variety of organizations, whose contributions are described in the acknowledgement section. The following sections cover a number of topics, related to the IIR clocks. These topics are: 1) Frequency Variation, 2) Performance of the IIR clocks especially as compared with the previous generations of GPS clocks, 3) Life Expectancy, 4) Drift Characteristics, and 5) Anomalies observed in the clock output.

FREQUENCY VARIATION

Cesium atomic standards, as a general rule, remain close to the ideal cesium clock frequency and have almost zero frequency drift. In contrast, rubidium atomic standards tend to have significant drift and cover a wide spread in frequency relative to the ideal rubidium frequency. There is little published data concerning the frequency spread of space quality rubidium atomic clocks. Therefore, the frequency distribution of the Block IIR RAFS clocks is examined.

The first issue that was encountered is that, because of their drift, it is hard to define a nominal frequency of a rubidium standard. This issue is addressed by arbitrarily selecting the nominal frequency to be the value in the ITT RAFS specification, given to Perkin-Elmer (formerly EG&G) in 1990. (Equation 1 is not the actual frequency, but the 510th sub-harmonic of the RAFS hyperfine frequency in the vicinity of 6.834 GHz.) All clocks will have a different frequency in a GPS orbit compared to the frequency of an identical clock on the earth, because of relativity effects. The relativity correction for the GPS orbit is given in (2). Combining (1) and (2) leads to (3) for the nominal value of the RAFS in GPS orbit.

$$\text{Rb on Earth frequency} = 13\,401\,343.9300 \text{ Hz} \tag{1}$$

$$\text{Relativity correction GPS orbit} = 4.467 \cdot 10^{-10} \tag{2}$$

$$\text{Rb frequency in orbit} = 13\,401\,343.9360 \text{ Hz} \tag{3}$$

The RAFS frequency values that were measured at the factory acceptance covered a normalized frequency range of

$$\delta f / f = -1000 \cdot 10^{-12} \text{ to } +650 \cdot 10^{-12} \tag{4}$$

The RAFS frequencies vary from the time of their acceptance at the factory to their values in orbit. Each of the clocks has its unique history due to a variety of conditions experienced between acceptance and in-orbit operation. These conditions include: 1) being power cycled many times, 2) subjected to various test procedures in and out of vacuum, 3) stored for a considerable time, and 4) being subjected to the rigors of transportation among various locations, including a launch to GPS orbit. We wondered whether the different stresses placed on the clocks might cause their frequencies to vary. Therefore, we examined the six clocks launched before 2003 and chose three frequency points, namely; a) the factory Acceptance Test (ATP) value corrected to correspond to GPS orbit operation, b) the initial frequency value when it was first tracked by National Imaging and Mapping Agency (NIMA) [2], and c) the value as of 2002-11-01. These three values are plotted in Figure 1 for the six satellites with Space Vehicle Number (SVN) values: 41, 43, 44, 46, 51, and 54.

Figure 1 shows the in-orbit active RAFS units have a fairly small spread (actually under $40 \cdot 10^{-12}$) except for SVN41, which has a very large variation of $362 \cdot 10^{-12}$ between the factory ATP value and its initial in-orbit value. All the changes in frequency, for the various RAFS from their initial values to the values on 2002-11-01, were negative and fairly small, with the largest magnitude under $50 \cdot 10^{-12}$. Our estimation is that if one includes the effects of drift for up to 10 years of operation, the frequency change during in-orbit operation will be less than $200 \cdot 10^{-12}$.

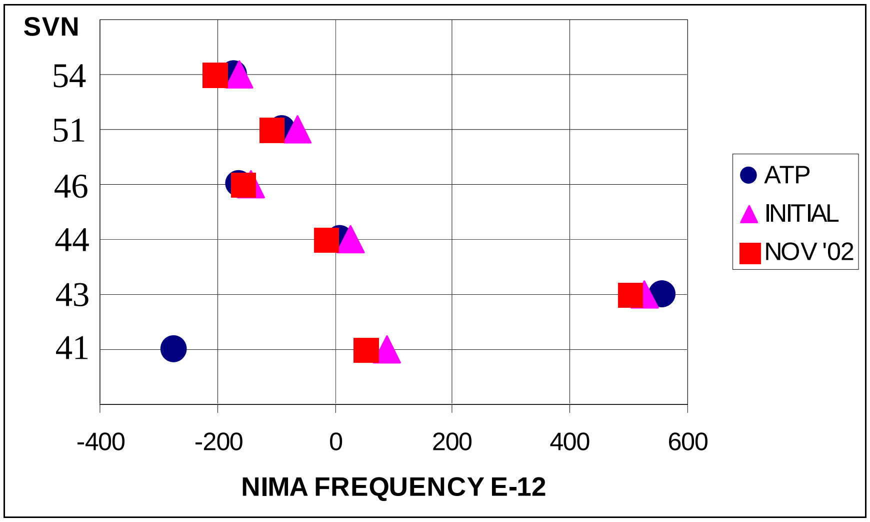

Figure 1. RAFS frequency history.

We wondered why there was such a large difference between the ATP frequency and the initial in-orbit frequency of SVN41 when the performance of SVN41 is similar to that of the other IIR in-orbit clocks. We tried to determine whether the difference was due to a mistake in the clock identity or that some severe shock to the clock might have accounted for the difference. A review of the various records did not reveal any reason for the difference. Thus, the cause of the difference remains unknown. The actual frequency value is not important, since the Time Keeping System (TKS) **[3]** is set up to cancel a constant frequency offset. However, it is worrisome that what we think is a rugged, high-precision clock has such a large unexplained frequency change between manufacture and in-orbit operation. Our intention is to continue to observe the IIR clocks and see if any other clocks have a large difference and if we can understand the observed behavior.

Performance and Stability

A key indicator is Estimated Range Deviation (ERD) performance. ERD is the difference between the current Master Control Station (MCS) Kalman Filter state estimate for the satellite's position (and clock) and the satellite's position (and clock) as broadcast in the navigation data along the path between the receiver and the satellite.

Figure 2 shows a typical sample of peak ERD ranking. Lockheed Martin, the Block IIR prime contractor, produces, on a weekly basis, an ERD rating for all the GPS satellites. The peak ERD is the largest magnitude of the observed 15-minute ERD measurements per satellite per day. These per-day peak values for each satellite are then averaged over a 7-day period to form a short-term ERD measurement and over a 90-day period to form a long-term measurement. The 7-day average is on the left and the 90-day average is on the right, for each satellite. The ranking of the satellites is done using the longer-term average. Figure 2 shows five of the first six satellites with the lowest peak ERD, namely SVN41, 43, 46, 51, and 54, are IIR SVNs.

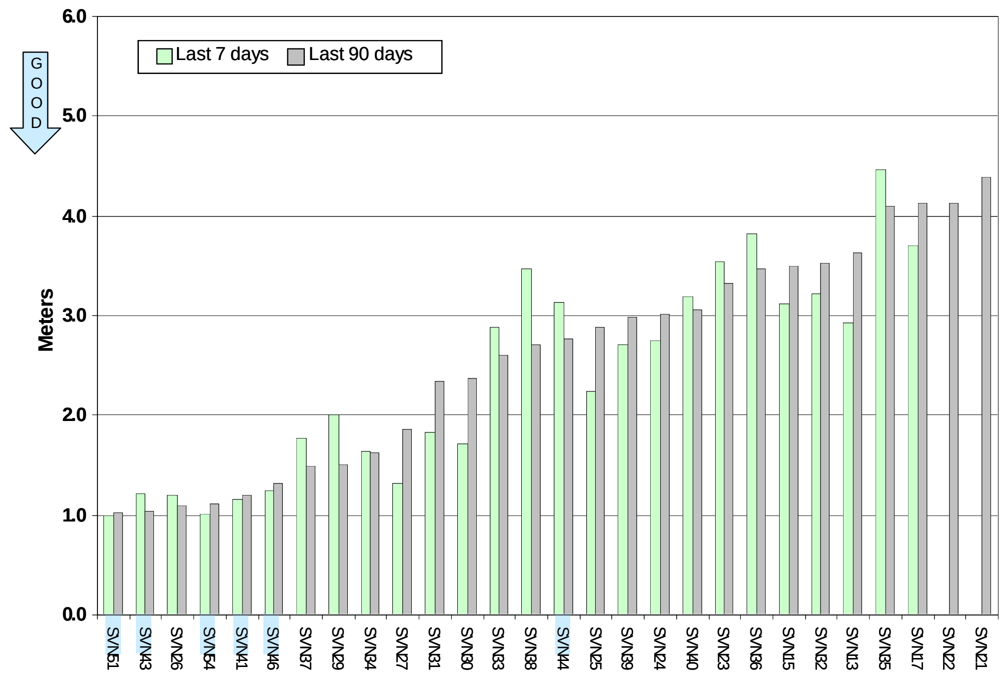

Figure 2. Ranking by ERD.

The lowest performance IIR RAFS clock (SVN44) is adequate and close to the middle ERD ranking. The remaining satellites, shown in this figure, are from GPS Blocks II and IIA, with 10 of them using rubidium clocks and 12 of them using cesium clocks. An examination of the positions of II/IIA cesium clocks and rubidium clocks shows that the IIR RAFS, as a family, are better than the IIA rubidium clocks, which are better than the II/IIA cesium clocks, as a family. There are exceptions to this general rule. In any case, one sees a general picture that the IIR satellites and clocks generally have better ERDs than the other satellites.

The Hadamard deviation is a measure of clock stability performance. The traditional measure of clock stability performance is Allan deviation. However, the current consensus in the community is that the Hadamard deviation is the preferred performance measure for the GPS rubidium clocks because the Hadamard deviation is: 1) unaffected by a constant drift, 2) converges in the presence of flicker walk and random run FM, 3) matches the three-state model used by the MCS, which compensates for these items, and 4) provides a good predictor for GPS rubidium performance **[4]**; see End Note.

The Hadamard deviation is a function of the delay, TAU, between three samples of the clock frequency. For the GPS constellation, which typically makes a prediction of the clock behavior, once a day, the Hadamard deviation, with a TAU of 1 day, is a good predictor of clock performance. Lockheed Martin generates a report of the behavior of the Block IIR satellites, on a quarterly basis. Figure 3 contains the chart of the SVN Hadamard deviations, for TAU of 1 day, for all the active GPS satellites, measured over the period of the third quarter of 2002. The satellites are ranked by the value of the Hadamard deviation, with the lower value having the better performance. Five of the first six satellites on the chart are IIR satellites and, as before, SVN44 is worse than the other IIR satellites, with a performance that is close to the constellation average. Also, the IIR clocks are generally better than the II/IIA rubidium clocks and the II/IIA rubidium clocks are generally better than the II/IIA cesium clocks.

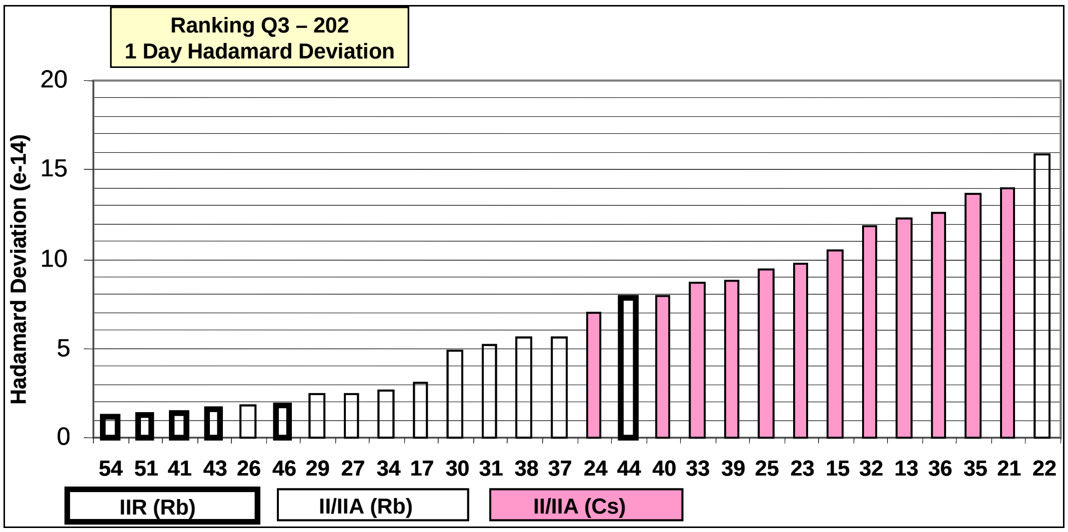

Figure 3. Ranking by Hadamard deviation.

Another indicator of clock performance is the User Range Error (URE) prediction vs. age of upload. Boeing generates a report of GPS performance on a monthly basis. Figure 4 describes the URE, which is an rms measure of the range error observed by a user, for the given month, as a function of the age of the upload. This curve is a hypothetical curve, because it describes the error that a user would see, if the upload, being used by the satellite, was made N days ago, where N ranges from 0 to 15 days. In actual practice, each GPS satellite receives an upload at least once a day. It is extremely unlikely that the active upload in a satellite will be over 2 days old. This curve allows one to see how well one can predict the future behavior of the clocks. There are three sub-figures in this chart, with one chart for Block II satellites, one chart for Block IIA satellites, and one chart for Block IIR satellites. The chart of URE vs. Age of Upload for Block IIR is far superior to the other blocks of satellites.

In summary, one can see that various measures for clock performance all point to the superior performance of the Block IIR RAFS, relative to the other clocks in the current constellation.

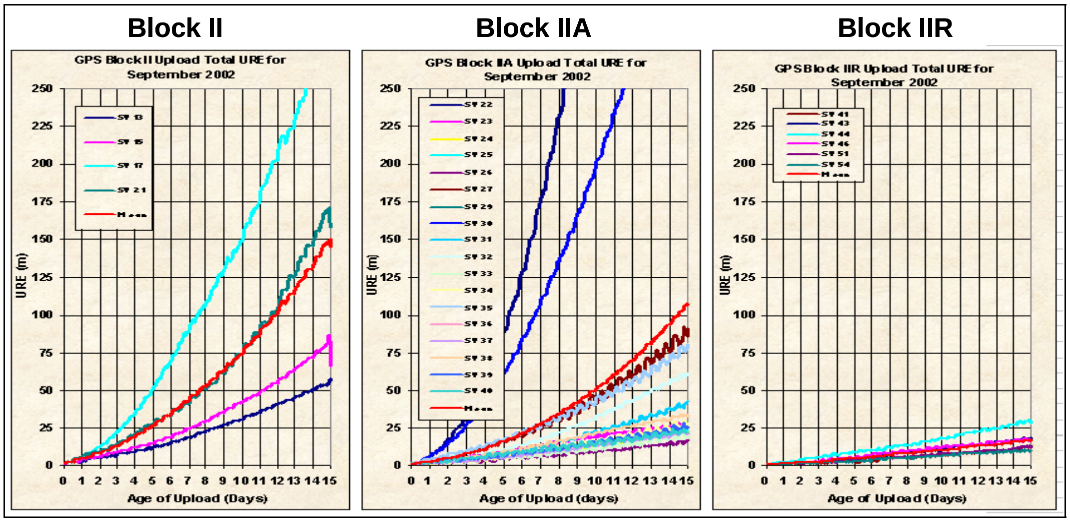

Figure 4. Comparison by age of data.

LIFE EXPECTANCY

If one examines the clock literature, there is much attention paid to clock performance with little attention to clock life. As far as the user is concerned, long life is as essential as good performance. This is well illustrated by the II/IIA rubidium clocks. Each II/IIA satellite has four atomic standards of which two are rubidium and two are cesium clocks. The II/IIA rubidium clocks perform better than the II/IIA cesium clocks. The life expectancy of the II/IIA rubidium clocks (~ 2 years) is worse than the life expectancy of the II/IIA cesium clocks (~ 6 years). The II/IIA satellites with an average service life of 10 to 12 years cannot survive over their lifetime using just their rubidium clocks and have used their cesium clocks for the majority of their service time. Some of the cesium clocks have very long lives. A classic example is SVN13, with 14 years of service life, the longest-lived GPS satellite to date. SVN13 has used a single cesium clock, which is safely within spec. SVN13, over its lifetime, has acquired a number of vehicle deficiencies, none of which have to do with the clock. SVN13 may possibly be retired in the next few years with two unused rubidium clocks, whose performances statistically are better than the active cesium clock. The obvious conclusion of this discussion is that we have to be concerned with long life as well as maximum quality performance from the GPS satellites.

There was significant concern, in some quarters, about the clock life expectancy issue in the early life of the Block IIR satellites. These satellites have three rubidium clocks and no cesium clocks. If the RAFS life were only 2 years (similar to the II/IIA rubidium clock), then the Block IIR satellite would be limited by clocks alone to an average of 6 years, which is less than the specified vehicle life requirement of 7.5 years. If the initial 7.5-year estimate for the average life of each RAFS is correct, then the three onboard clocks have enough life for much more than the specified vehicle life. Thus, there has been much interest in the life of the RAFS. Conceptually, one would like to have a full life test of these clocks, in which many clocks are operated on a test until they fail or are well beyond the projected 7.5-year vehicle life requirement. After the life test was run, one could then start to use these clocks. This approach was rejected because: 1) it would take too long, 2) it would be too costly, and 3) the GPS project office was not willing to delay the use of the RAFS until the end of a life test. Instead, a life test was started at NRL with two flight-qualified RAFS clocks a few months before the first projected Block IIR launch **[5]**. The life test data augments the data gathered from the operation of the in-orbit RAFS clocks.

NRL provided life data for all the GPS clocks, starting with the first GPS satellite through November 2002 **[6]**. When the data were examined, it was not clear how the life of the clocks should be defined, such as: when the clock was on, even in the manufacturing facility, only after delivery to the customer, only in-orbit operation, or only when the clock was healthy and delivering GPS service. Although one could justify various options, Lockheed Martin defines Block IIR life as the time the clock is ON while it is in orbit, whether it is selected or not. For the non-Block IIR clocks, the life was defined as all the time the clock was ON and selected to be the operational clock, because that was the way the data were presented by NRL. In practice, the estimates of GPS clock life obtained by the Lockheed Martin and the NRL techniques do not differ significantly.

From an analytical point of view, one would like to have the distribution of the life of a clock independent of other items. The measurements present an issue of unexpired life. This issue is that data stop when the satellite is taken out of service with a still-functioning clock and it is therefore not clear when the clock would have failed. This issue shows up in two forms: a) the clock on the in-orbit satellite is still functioning well but some other critical component fails, which causes the GPS satellite to be retired including the clock (it is not determined how long the clock would have lived), and b) the clock and the satellite are still alive and functioning at the present time and it is not known how much longer it will operate.

One approach to this issue is to just count the data we have. This approach tends to give a smaller estimate of the clock life than the correct value because unexpired life is not counted that should be counted to get a true estimate of the average clock life. An alternate approach is to ignore the clocks that have unexpired life and just count the clocks for which data exist from the time they were turned on in-orbit until they die. This approach, unfortunately, also gives a smaller estimate of clock life, because it counts those clocks that failed early and tends to ignore the contributions of the longer-lived clocks. One can illustrate a problem with the service time approach by assuming an artificial case where the cesium clocks live exactly 9 years, but the satellite typically dies at 12 years because the reaction wheels used in the attitude control system die at 12 years. In this case, the first cesium would last 9 years and the second cesium with 9 years of life will be turned off after its third year because the SV wheels have failed. In this case, one can say that the two cesium clocks gave 12 years of service and so the average life of a cesium clock is 6 years, which clearly is in conflict with our assumption of a 9-year life. A possible approach to this issue is to get some estimate of the unexpired life distribution, assuming some model of the clock life, and add this estimate to the clocks with unexpired life.

The life of the six active in-orbit Block IIR rubidium clocks and the life test clocks as of the end of 2002 were examined. There are two other Block IIR clocks that have been turned on for a short time, which were not considered. The first such clock is clock 3 of SVN43 that was powered for 10 months, at the beginning of the life of SVN43, as a test to see how the Block IIR system would behave when there were two clocks active. Once it was determined that the satellite worked well with two powered clocks, clock 3 of SVN43 was un-powered in order to extend the life of this clock and reduce the load on the power system. This clock was not considered in our life analysis, as there was little data about its longevity. A second clock that was powered for a short time was clock 1 of SVN44. This clock was the first clock turned on for SVN44 and it exhibited excellent performance. However, when it was powered, the heater system that stabilizes the base temperature of this clock did not function. Although, this clock exhibited excellent performance, the Control Segment (CS) personnel were concerned that its performance might degrade severely during the eclipse season as the satellite passes in and out of sunlight. The baseplate would have temperature variations that could affect the clock frequency. Some felt that clock 1 would function acceptably even without the baseplate heater, but the decision was made to go to clock 2 of SVN44. Life data on clock 1 were ignored because the clock was still functioning and the issue was the possible performance degradation during eclipse season. The data listing also ignores the clocks that have been launched since 2003, because they all are operating properly but their life is too short to be significant for an estimate of life expectancy.

The basic conclusion from the existing Block IIR life data is that, as of the present, there are eight clocks, with two over 6 years, one over 5.5 years, and the shortest over 2 years, and that all were still operating. Because all the clocks are still operating, the life expectancy of the clocks cannot be calculated. However, a reasonable estimate is that the average clock life is greater than 6 years. When the life of the pre-Block IIR rubidium clocks was reviewed, a much different picture is evident. The average service life of these clocks was close to 2 years and even if one added some estimate for the unexpired life, the average life was probably not over 3 years. The Block I and Block IIA rubidium clocks had significant infant mortality, defined as clock failure within 0.5 year of power on. Twenty-three Block I rubidium clocks had failed in service, while three were still functioning when the satellite was retired due to the failure of some other component. Of the total of 26 Block 1 rubidium clocks that were powered, nine failed in the first 0.5 year. Twenty-eight Block II/IIA rubidium clocks were powered. Eighteen clocks had failed, 10 had survived in service, and seven had failed in the first half-year **[7]**.

An obvious conclusion was that the Block I and II/IIA rubidium clocks had a significant infant mortality, but the Block IIR RAFS has an insignificant infant mortality. Thus, the Block IIR RAFS has a life expectancy that is much different than that of the preceding rubidium clocks in two ways: a longer average life span and a lower infant mortality. Unfortunately, an accurate estimate of RAFS infant mortality does not exist, because none of these clocks has exhibited this phenomenon. Currently, the demonstrated average life expectancy is over 5 years (and increasing), but how long it might be is unknown, since none of the Block IIR clocks has failed. However, the satellite life is liable to be limited to a value near 12 years, because of the wear-out of the wheels in the attitude control system and/or the degradation of other components. With the clock life average of over 5 or 6 years, the average life offered by three clocks becomes over 15 to 18 years. This means that the clocks, as a group, are likely to outlive the other critical components and, thus, the clocks are not going to be a significant cause of GPS satellite failure.

To put this matter into perspective, the following items relative to the GPS failure mechanisms are noted. The first four GPS satellites failed due to clock loss, so that, initially, clocks were the predominant failure mechanism. Since then there have been 14 GPS satellite failures, with the following breakdown: a) 6 wheels, b) 4 clocks, c) 2 power, d) 1 TT&C (Telemetry and Control), e) 1 Navigation Data Unit (NDU). Our estimate is that the Block IIR satellites will not fail because of clock failure. Currently, three satellites (all of them Block II/IIA) are on their last clock. All of these satellites are over 9.5 years old, so that even if the satellite fails due to clocks, there is little life left in the vehicle, because of the impending end-of-life failure of other components.

There are significant issues that need to be followed with respect to the clock life. The primary item is, of course, to monitor clock life and get a better estimate of clock life as data are accumulated. Secondly, the Block IIR clocks will have significant shelf life. In the worst case, a clock could have been manufactured in 1994 and may not get launched until the last Block IIR vehicle, which may be in the vicinity of 2010, so there may be a ground shelf life of 16 years. If the clock is the last clock to be activated on a satellite that lasts in the vicinity of 10 to 15 years, there could be an additional 10 to 12 years, on storage in orbit, with the more severe space-based environment than the ground environment. Current experience and measurements indicate that this storage life will not seriously impair clock life. However, this is a matter that should be tracked. Thirdly, rubidium clock performance behavior may change significantly with life without any early warning of clock failure. Most cesium clocks do have an early warning in terms of an increase over several weeks of the beam voltage. Thus, there should be an effort to track the Block IIR RAFS operation over their life to gather data concerning these various issues.

If, as our data indicate, clock life is not a significant cause of Block IIR satellite failure, then strategies to improve the performance of the satellite broadcast signals should be considered. For example, the best clock should be selected initially, so that the unused life should be from the worst performing clock and not the best performing clock. Also, if a clock and its behavior turn out to be sub par, for whatever reason, one should consider changing the selection to a better clock, even if this means deselecting a clock while it still has significant life left. If the second clock fails, one can always go back to the original clock. A desirable feature of the Block IIR rubidium clocks is that they can be operated continuously, with no need to stop SV service for clock maintenance. In fact, of the six Block IIR clocks that were in service for the last half of 2001, five provided 100% availability in this period and the remaining satellite (namely SVN51) had 100% availability except for one 7.5-hour scheduled outage to perform a station-keeping maneuver.

DRIFT CHARACTERISTICS

Frequency drift is a key characteristic of the Block IIR rubidium clocks. Absolute phase is a quantity, which is initialized from an outside reference and not related to the clock characteristics. Also, the frequency of the Block IIR clocks varies widely without seeming to affect clock performance. The drift characteristics of the Block IIR clocks seem to be intrinsic to rubidium clocks. In practice, it is difficult to directly measure the clock drift. Instead, one typically measures the phase or frequency of the clock directly and then performs a difference or derivative operation to measure the drift. The difference operation implicitly contains a filter. The effect of the filter is very significant and the same data can appear very different, to the eye, depending on the filter. Figures 5 and 6 illustrate this effect. Figure 5 shows the drift for SVN 41 clock for the year 2001 with a wideband filter and Figure 6 shows essentially the same drift over roughly the same period with a narrowband filter and coarser sampling. These plots of the drift are used for different purposes. The narrower filter presentation makes it is easier to see long-term drift or frequency steps. The wider filter is useful in detecting fast changes that are suppressed by the narrower filter.

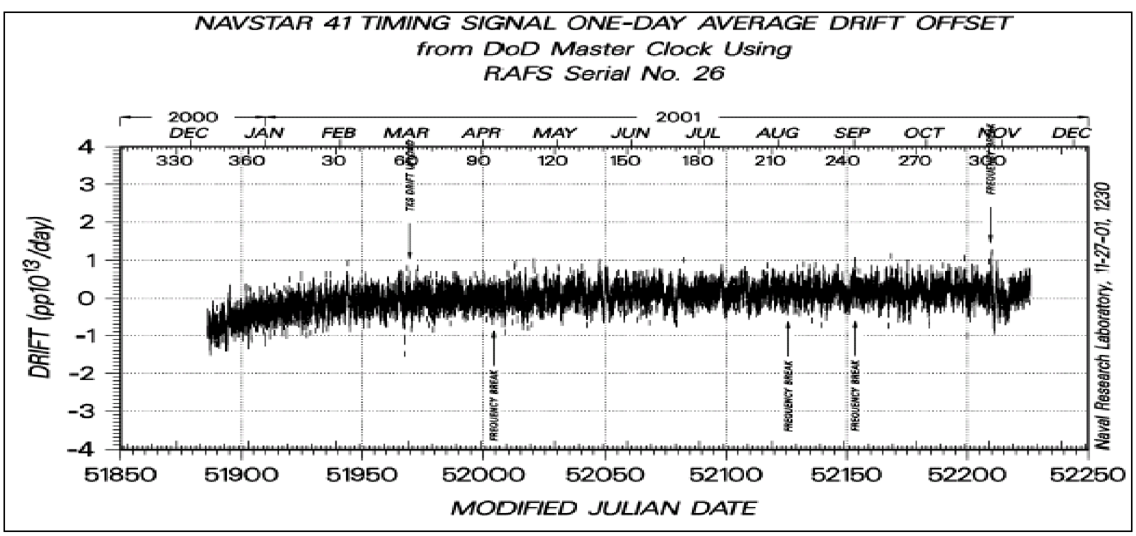

Figure 5. Drift with wideband filter.

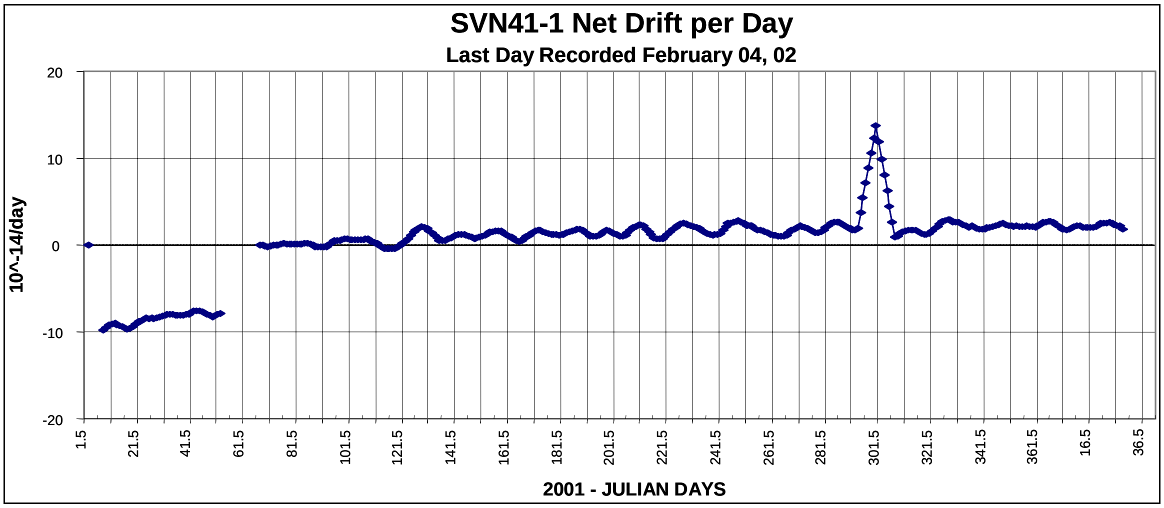

Figure 6. Drift with narrowband filter.

The drift characteristics of the various Block IIR clocks follow a common pattern vs. time. This is illustrated in Figure 7, which contains a narrowband filter plot of the drift, for three Block IIR clocks vs. time. The three drift curves have been shifted in time so that the average drift for each of these three clocks has a value of $-10 \cdot 10^{-14}$/day, when the time axis is 0 days. All three drift curves start with a very large magnitude negative drift. These curves have a generally positive slope and get closer to zero over time but never reach an average value of zero drift. Because the average slope of the drift curve is always positive, the curves have negative drift values, with a magnitude of over $10 \cdot 10^{-14}$/day for negative values of time and values with a magnitude of less than $10 \cdot 10^{-14}$/day for positive time. The decay of the drift is initially rapid and then diminishes with time. All the curves get closer to zero drift than $5 \cdot 10^{-14}$/day and currently all the six Block IIR clocks with over 2 years life are between $-4 \cdot 10^{-14}$/day and zero. The SVN46 clock clearly decayed the fastest initially and, as of the present, is the closest to zero drift, with an average drift of about $-1.8 \cdot 10^{-14}$/day. This drift value is remarkably low for a rubidium clock.

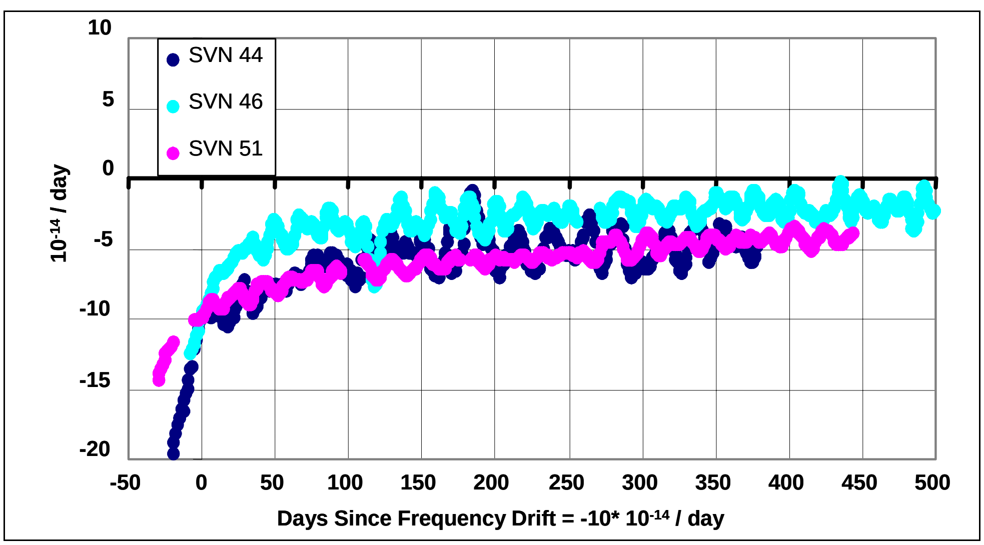

Figure 7. Various RAFS drift characteristics.

A typical characteristic of the Block IIR clocks is that the drift seems to initially come down rapidly with a fast time constant and then switches to a slower time constant. This characteristic is illustrated in the SVN46 plot. One hypothesis is that the different time constants are associated with different processes inside the rubidium clock. According to this hypothesis, the faster time constant process dominates the drift behavior initially, and then the effect of the fast time constant process becomes much smaller than that of the slower time constant process, which dominates after the crossover. As seen in Figure 7, SVN46 has the lowest drift value and the fastest time constant, while SVN51 has the largest magnitude drift value and the longest time constant of the three clocks shown in this figure.

The Time Keeping System (TKS) of the Block IIR satellite allows the CS to cancel out the clock drift. This capability is used to control the frequency and phase of the satellite broadcast clock, with nominal frequency 10.23 MHz. The RAFS is an open-loop oscillator. Its output frequency is in the vicinity of 13.4 MHz and is unaffected by the TKS system. However, the satellite output signal is locked to the RAFS with phase, frequency, and drift offsets, so that the drift in the satellite broadcast signal differs from the RAFS drift by the drift correction, which can be the same value or smaller or larger than the RAFS drift. Thus, while the RAFS drift is always negative, the broadcast drift can be zero, negative, or positive.

One can see the effect of drift corrections on the SVN44 RAFS in Figure 8. The SVN44 clock started with a positive frequency of about $+16 \cdot 10^{-12}$ and then decayed to $-1.25 \cdot 10^{-12}$, at which point a drift correction was made to cancel out most of the drift. When the frequency reached $-3.45 \cdot 10^{-12}$, the drift correction was increased so that the net drift was positive and the frequency had a positive slope. The frequency then rose so that it reached zero and finally reached the value of $1.211 \cdot 10^{-12}$ on 2002-11-11, which is the last point plotted in Figure 8. This process of phase and frequency management has kept the frequency of the Block IIR clocks between $\pm 4 \cdot 10^{-12}$ after an initial transient period. Note that one of the major reasons for using frequency drift corrections is that these corrections start off as subtle corrections that don't cause any jumps in phase or frequency. Thus, frequency drift correction can be performed while the satellite is healthy, with no interruption in service and without causing any significant degradation in the GPS performance.

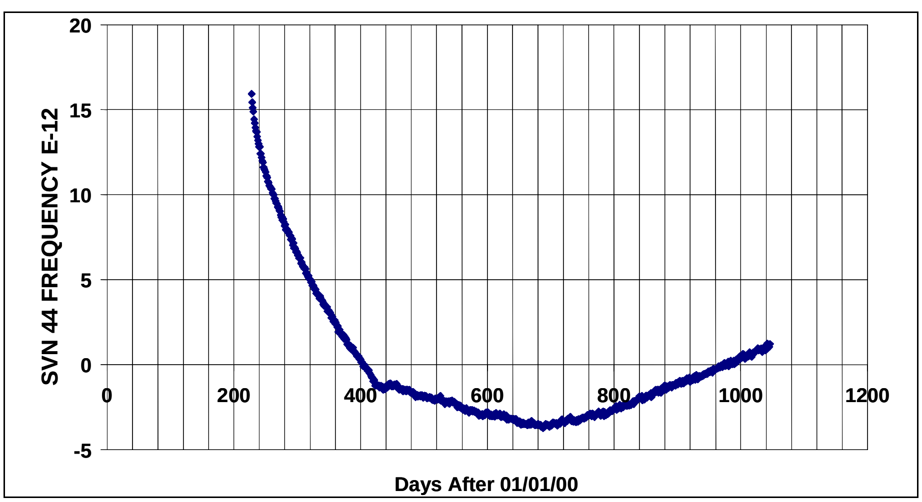

Figure 8. SVN44 frequency adjustments from 01/01/00 to 11/11/02.

Figure 9 shows histograms of the frequency error distribution of the GPS clocks. The IIR rubidium clocks have a frequency distribution much narrower than the cesium clocks, which shows the effectiveness of the phase and frequency management technique. The effort to manage the phase and frequency distribution is minimal. The original six Block IIR RAFS will require only one frequency drift correction in 2003. The effectiveness of the frequency drift corrections depends on being able to predict the frequency drift behavior of the RAFS over a number of years. This is not simple to do, as there is some unpredictable variability in the RAFS drift.

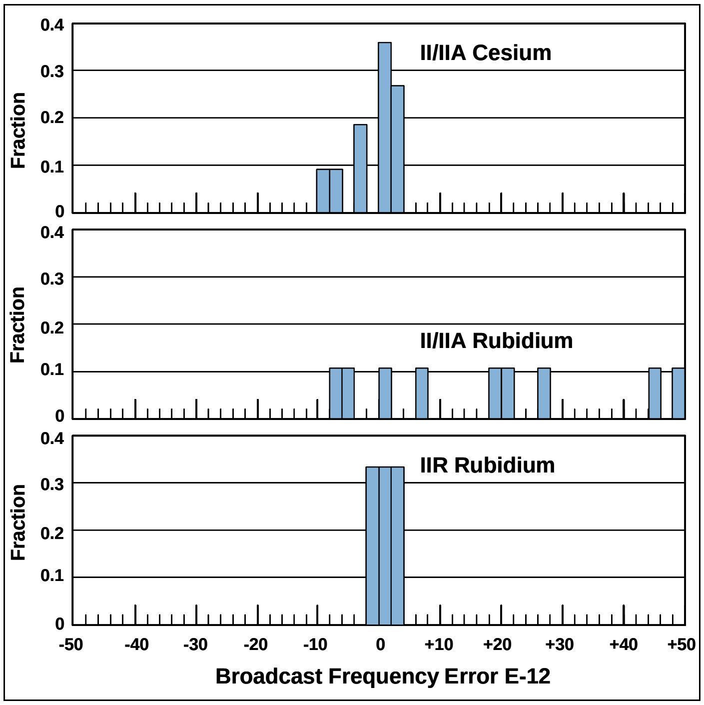

Figure 9. Frequency distribution histograms for GPS clock families.

The RAFS drift has some interesting phenomena. In particular, there seems to be spectral peaks associated with many of the RAFS. One can see this phenomenon in Figure 10, which shows an oscillation with a period on the order of 10 days in the SVN46 clock. Similar spectral peaks with various periods can be seen in other RAFS.

In conclusion, we see a number of interesting characteristics associated with the Block IIR RAFS. We have a useful technique for minimizing net drift.

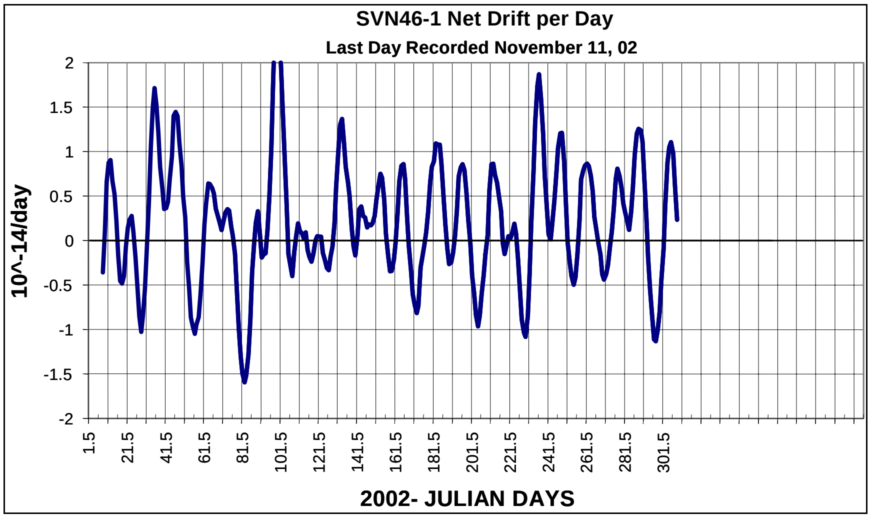

Figure 10. SVN41 frequency for year 2001.

ANOMALIES

The Block IIR RAFS is the victim of its own success. The RAFS is the newest clock in the GPS constellation, with the best performance and the cleanest signal. This situation has encouraged the clock observers to search more carefully for abnormal behavior. As a result, there is a long list of anomalies that are associated with the RAFS in particular and the broadcast signal in general. Many of these anomaly types are present in pre-Block IIR clocks, but because these clocks were noisy, the observers of the clocks did not or could not distinguish these anomalies within the noise. When the same anomaly types appeared in the Block IIR RAFS, they were easier to detect. The list of anomalies includes: a) upload command effects, b) relativistic frequency effects, c) phase excursions, d) phase steps, e) frequency steps, f) drift steps, g) spectral features, and h) unexpected oscillations.

The broadcast signal reveals the effects of upload commands, which can change the broadcast phase, frequency, and drift. The present goal is to perform a phase and a frequency adjustment at clock startup, before the clock is declared operational and to limit the changes to frequency drift commands once the clock is operational. If special circumstances, such as a faulty upload, demand it, the required phase and frequency commands can be issued. As long as the observer understands that the observed effects are the results of a conscious command, these anomalies present no problem. Because the Block IIR clocks are so clean, a relativistic frequency step effect associated with minute orbit changes was observed **[8]**. Once the source of this effect was recognized, there was no problem about arranging to compensate for this effect that affects all clocks, according to the standard general relativity theory. Thus, there is a simple approach to handling the effects of upload commands and relativity.

The phase excursion effect is a frequency step, not in the RAFS, but in the VCXO, that generates the broadcast signal in the TKS. The phase excursion is called by this name because the VCXO frequency step initially causes a phase excursion, which is quickly corrected by the TKS phase-locked loop, so that there is no long-term change in phase due to a VCXO frequency step. The phase excursion initially could cause the broadcast code to go to Non-Standard-Codes (NSC), which causes an unscheduled loss of service to the users. This effect was detected and a software change was implemented **[9]** to prevent the loss of service. The TKS software was modified to more properly monitor such effects to better determine the presence of such frequency steps in the VCXO.

Phase, frequency, and drift steps have all been detected in the RAFS. The phase steps have been very rare. The key item is the presence of frequency steps, which are a well-known clock phenomenon. Frequency steps have been observed in the RAFS life test and carefully documented. Figure 11 shows a frequency step that occurred in orbit on 2002-10-22 (DOY 301). This frequency step did not cause a loss of service, but it did cause a number of extra uploads and higher ERD values until the Kalman filter, which tracks the clock, could adjust to the new frequency. The frequency step, shown in Figure 11, is a step in frequency that leaves a permanent effect. On the other hand, we have frequency steps in SVN54, as shown in Figure 12, which shows a frequency step that is triangular in shape as the frequency changes and then returns to the previous frequency curve (days 23 and 54). This class of frequency steps has not only a frequency step, but also a drift step characteristic.

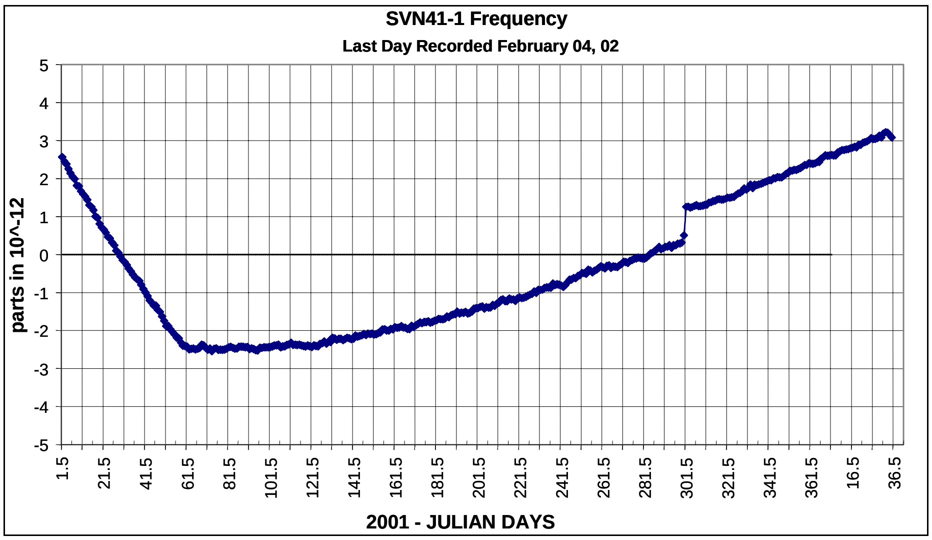

Figure 11. SVN41 frequency for year 2001.

The effect of the frequency steps is significant and it shows up as an increase in the Hadamard deviation associated with the clock. If one could remove the frequency steps, then the measured Hadamard deviation would improve significantly. This is illustrated by Figure 13, which shows the Hadamard deviation with the actual frequency steps contained in the data (upper curve) and with the frequency steps removed (lower curve). This difference is important, as it affects the 1-day Hadamard value, which is related to the average ERD. Some have argued that the only real clock measure is the one produced by the RAFS and there is no validity to the Hadamard deviation with the frequency steps removed. As an example, assume that a sharp frequency step occurred and the user is a surveyor, who uses National Imaging and Mapping Agency (NIMA) postprocessed data to produce his results. If NIMA examines the clock data, after the fact, detects the frequency step, and totally removes it from the postprocessed data used by the surveyor, then the surveyor's output is unaffected by the frequency step and his Hadamard deviation is effectively the plot with the frequency step removed. If one wants to examine the effect of a frequency step on a real-time user, then one gets into questions of how fast the frequency step is detected, how much phase error accumulates statistically before there is a correction for the frequency error, and what is the effect on the real-time user in this environment with frequency steps and corrective measures. This is an interesting area that is worthy of study.

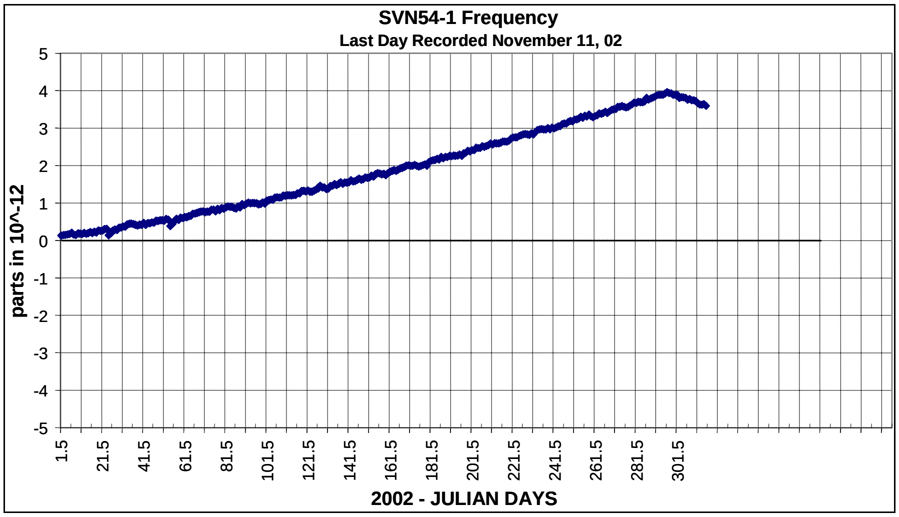

Figure 12. SVN54 frequency for year 2002.

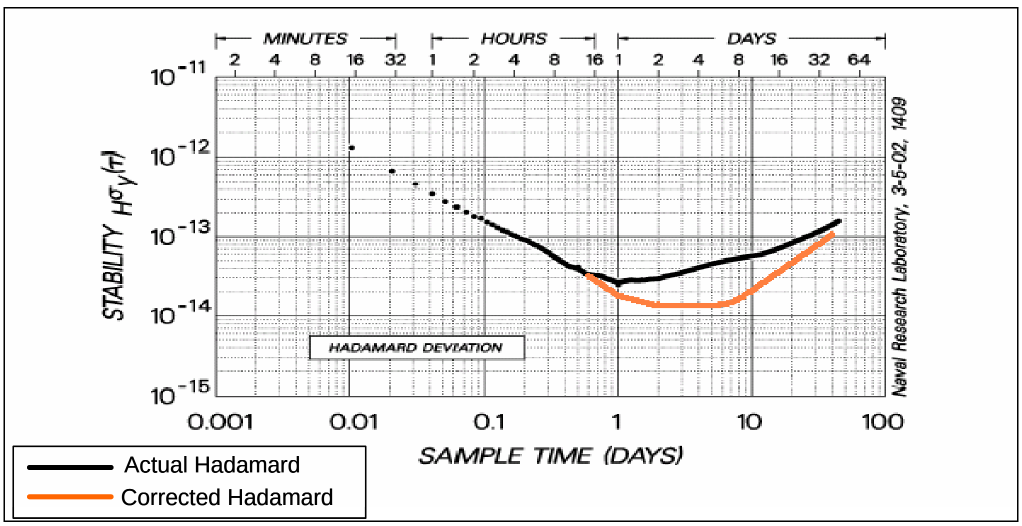

Figure 13. Hadamard deviation without frequency steps (lower curve) and with frequency steps included.

There have been a number of spectral features detected in the Block IIR clocks. The most significant item is a spectral peak with a period of about 2.5 days in clock 2 of SVN44. This spectral peak is responsible for the fact that SVN44 is out of family compared to the performance of the other Block IIR clocks. A number of solutions have been considered for this clock. However, as of the present, this clock is still the active clock in SVN44 with no major change in its performance. Unexpected oscillations were detected in the TKS loop under certain conditions. A software upgrade to the TKS control loop was made to remove any tendency for oscillation. The ability to reprogram the satellite is an essential part in dealing with these issues

CONCLUSIONS

There are a variety of conclusions resulting from of this survey of the Block IIR clocks. These clocks are clearly the best in the current constellation. Although we did not discuss this topic in the previous sections, we assume that the next generation of rubidium clocks will have even better performance **[1]**. The expected life of the Block IIR rubidium clocks is distinctly better than that of the previous generations of rubidium clocks. The Block IIR satellites lives will not likely be limited by clock failures. There is currently a limited observation of the lives of these clocks, which are in their adolescent lives. Continued observation of the clocks is needed in order to obtain an understanding of all parts of the lives of the clock and, in particular, the later portion of their lives. There were several issues raised in this survey that remain to be addressed. The frequency steps have a significant effect on the clock performance. There have been some considerations made on how one can detect and mitigate the effects of the steps in addition to the efforts to improve the clock design to reduce or eliminate the occurrence of these steps. The length of time from the start of this design in 1980 to the time when the last Block IIR clock will be retired in the 2020s is an extraordinary length of time. This raises the question of how we, as a technical society, can capture the lessons to be learned from this design and make sure they are not lost in future clock designs.

ACKNOWLEDGMENTS

This overview would not be possible without the wholehearted assistance of many organizations and individuals within these organizations. The following is a list of contributors: Perkin Elmer (formerly EG&G), William Riley (now at Symmetricom) and John Vaccaro, are responsible for the design and manufacture of the Block IIR RAFS. We are grateful to the US Air Force and, in particular, 2SOPS 1st Lt. Stephen Bolt and Steven Hutsell, who control and manage the IIR clocks, for their cooperation in our work with these clocks. Steven Hutsell also reviewed the draft paper and made a number of suggestions, which we incorporated. Lockheed Martin is the prime contractor for the Block IIR satellites and provided us with their weekly and quarterly reports as well as with assistance on many other items. The individuals involved with this assistance include Karen Gauthier and Williard Marquis, Jr. Boeing, the manufacturer of the Block I and II/IIA GPS vehicles, monitors all the GPS satellites and provides valuable data about the Block IIR satellites in their weekly and monthly GPS reports. Jack Taylor is the generator of much of these data and has always been very helpful. Much of our work is based on data from the NIMA Web site. Our analysis is based on the NRL reports on both the clocks in orbit and those in the NRL life test. NRL also gave us the life histories of all the GPS clocks. The individuals associated with the NRL effort are James Buisson and Jay Oaks. Lastly, we want to acknowledge the assistance of the whole ITT team, and, in particular, Jay Phelan, Todd Dass, Jeffrey Harvey, Randell Reith, Harris Rawicz, John Petzinger, and John D'Agostino. We apologize to the many others who assisted us with material that was used in the paper, but were not included in this list because of space limitations.

REFERENCES

[1] T. Dass, G. Freed, J. Petzinger, J. Rajan, T. Lynch, and J. Vaccaro, 2003, *"GPS Clocks in Space: Current Performance and Plans for the Future,"* in Proceedings of the 34th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 3-5 December 2002, Reston, Virginia, USA (U.S. Naval Observatory, Washington, D.C.), pp. 175-192.

[2] NIMA Web site, *http:// 164.214.2.59/GandG/sathml/l*

[3] H. C. Rawicz, M. A. Epstein, and J. A. Rajan, 1993, *"The Time Keeping System for GPS Block IIR,"* in Proceedings of the 24th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 1-3 December 1992, McLean, Virginia, USA (NASA Conference Publication 3218), pp. 5-16.

[4] S. T. Hutsell, 1996, *"Relating the Hadamard Variance to MCS Kalman Filter Clock Estimation,"* in Proceedings of the 27th Annual Precise Time and Time Interval (PTTI) Applications and Planning Meeting, 29 November-1 December 1995, San Diego, California, USA (NASA Conference Publication 3334), pp. 291-301.

[5] *"NAVSTAR Analysis Updates and Life Test Reports,"* U.S. Naval Research Laboratory, Washington, D.C.

[6] J. Oaks and J. Buisson, *"Life of All GPS Clocks, Activated and Deactivated Lifetimes,"* private communication, e-mail from NRL to ITT, 6 November 2002.

[7] *"Almanac,"* GPS World Showcase, **GPS World**, December 2002, p. 56.

[8] M. Epstein, E. Stoll, and J. Fine, 2002, *"Observable Relativistic Frequency Steps Induced by On-orbit Changes,"* in Proceedings of the 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 27-29 November 2001, Long Beach, California, USA (U.S. Naval Observatory, Washington, D.C.), pp. 493-507.

[9] J. Petzinger, R. Reith, and T. Dass, 2003, *"Enhancement to the GPS Block IIR Timekeeping System,"* in Proceedings of the 34th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 3-5 December 2002, Reston, Virginia, USA (U.S. Naval Observatory, Washington, D.C.), pp. 89-107.

END NOTE

One can also justify the use of the Hadamard deviation over the Allan deviation as follows. Allan deviation is sensitive to a constant drift, while the Hadamard deviation is unaffected by a constant drift and instead is sensitive to changes in drift. We examined the RAFS frequency drift, after the RAFS had stabilized for over half a year. We looked at the range of TAU from 1 to 10 days, which is a significant range for the manner used to operate the GPS predictions. In these circumstances, the RAFS typical frequency drift is fairly constant over TAU days and its frequency drift variation is much smaller than the average frequency drift level. Thus, for TAU between 1 and 10 days, the Allan deviation is much larger than the Hadamard deviation. Since the GPS performance correlates well with the Hadamard deviation and not the Allan deviation, one can understand why various experts prefer to use the Hadamard deviation to evaluate the GPS rubidium performance. Note that it is not obvious that the RAFS frequency drift is fairly constant over a 1- to 10-day period. This is because one tends to look at the RAFS frequency drift over the years of clock life and sees a wide variation in frequency drift over its life, as much as an order of magnitude. On the other hand, one has to concentrate on the variation over a few days to notice that the frequency drift is fairly constant over this time span.

QUESTIONS AND ANSWERS

KEN JOHNSTON (U.S. Naval Observatory): Could you tell some of the longer-life clock technologies that would make the II-R's better than the previous rubidiums?

MARVIN EPSTEIN: I do not really know. A lot of this is buried in the lore of the clock manufacturers. I think there have been other long-life clocks that that have been rubidiums. I think there is something that happened with the 1 and 2-A rubidiums that somehow limited their life. I am not sure what you might do to make it work. I think in the 1 and 2-A systems, the cesium clocks were the top clocks. Those are the ones that they really tried to do. Lockheed sort of put a pile of components in as their own group, went over, and assembled the rubidiums.

The II-R rubidiums have a long history, as I've said. The guys that did it were nothing but clock people. They put their lives into it. They got a crew that worked on nothing but these clocks and somehow the difference shows up. To go over and point out the reasons, I don't really know. You would have to go and ask them. Somehow, it is a fantastic product.