ENHANCEMENTS TO THE GPS BLOCK IIR TIMEKEEPING SYSTEM

Mr. John Petzinger, Mr. Randall Reith, and Mr. Todd Dass
ITT Industries
Aerospace/Communications Division, NJ Technology Center
Clifton, NJ 07014-1993, USA
john.petzinger@itt.com
Tel: +1 858 578 3080 x210; Fax: +1 858 578 5371

Abstract

The ITT Industries-developed GPS IIR satellite payloads have been on-orbit since 1997, providing outstanding signal-in-space performance. Much of credit for this outstanding performance can be given to the GPS IIR Time Keeping System (TKS), and the GPS-IIR spacecraft bus, which keeps the payload in a mechanically and thermally stable condition. The TKS consists of a VCXO phase-locked to a* rubidium atomic frequency standard *(RAFS) via a control loop implemented in software. Because the TKS is constantly comparing the phase of a VCXO to the phase of the RAFS, it can autonomously protect GPS users from random frequency variations that might occur in either component, and also from potential system failures in the TKS. In addition, because the TKS control loop is implemented in software, it can be modified or tuned to handle these random events in an optimal fashion.

Although SVN 43's VCXO is atypical, we characterize the frequency of occurrence, magnitud, and seasonality of its frequency variations. We then describe how the addition of a nonlinear error processor to the TKS control loop, called an Adaptive Time Constant (ATC), was used to meet the challenge of the frequency variations. The rationale for ATC, an approximate maximum likelihood error detector, is presented. By analyzing on-orbit telemetry, we show how the TKS, after activation of ATC, has robustly shielded GPS users from VCXO disturbances that would have caused outages. Data from both SVN 43 and SVN 46 show that even during stable periods, ATC reduces undesirable noise from the TKS output. Using a detailed TKS simulation, we show that the TKS with ATC still protects users from potential anomalies in the RAFS. The conclusion is that ATC has improved TKS timing stability and TKS availability on the SVs that are using it.

1 EXECUTIVE SUMMARY AND HISTORY

The Time Keeping System (TKS) is the system in the GPS Block IIR satellite that provides an accurate 10.23 MHz signal to the code generators and to the L-Band subsystem of the navigation payload. It implements a *control loop* in software that continuously tunes a *voltage controlled crystal oscillator* (VCXO) so that it follows the timing stability of a *rubidium atomic frequency standard* (RAFS). Because two independent oscillators are compared every epoch, the TKS is able to detect anomalies in either the RAFS or in the VCXO and, if necessary, protect users from using a degraded signal by rapidly switching to non-standard codes (NSC) within seconds of the failure.

The Adaptive Time Constant mode (ATC) was added to the control loop software of the Block IIR Timekeeping System (TKS) as part of Operational Build (OB) 8, which was first uploaded to SVN 43 on 10/10/00 and then to SVN46. Since these uploads, ATC has improved the availability and the performance of the TKS on both SVN 43 and SVN 46. We will present the rationale for ATC and the evidence of improved availability and improved timing stability.

The impetus for developing the ATC came from the evidence that the VCXO in SVN 43 was experiencing periods of timing instability, particularly during eclipse periods. We will present the evidence for this instability in the hope that it may be useful for those characterizing and improving the performance of crystal clocks in space.

2 REVIEW OF TKS AND ATC FUNDAMENTALS

The purpose of the GPS Block II-R TKS shown in Figure 1 is to tune the 10.23 MHz digitally controlled VCXO to produce the GPS navigation signal with the timing accuracy of RAFS. By linking the VCXO to the RAFS using a control loop controlled by software, it is possible to precisely adjust the frequency and phase of the TKS output, to cancel drift of the RAFS once it has been characterized, and to detect any anomalous RAFS frequency or phase excursions. If a phase excursion larger than the *failure threshold* occurs twice in a row, the L-Band output is switched to non-standard codes autonomously.

2.1 TKS HARDWARE AND CALCULATION OF PHASE ERROR

Figure 1 shows that frequency divider circuitry produces a pulse once per epoch from the digitally controlled VCXO and once per epoch from the RAFS, which has a fixed frequency. A phase meter measures the phase relationship between the reference epoch from the RAFS divider chain and the system epoch from the VCXO divider chain. The phase meter uses a 600 MHz oscillator to provide measurements of this phase difference with an accuracy of ±1.67 ns, which is read by the software. The rest of the processing in the diagram is done in software run in the Mission Data Unit (MDU). The MDU software calculates a phase error by subtracting out the intentional variation in phase due to dither (if dither is enabled), and then subtracting the expected phase difference due to other factors. The phase error is scaled as a fraction of an epoch and is typically less than 2 ns or $1.3 \cdot 10^{-9}$ epochs. If the phase error is positive, it indicates that the system epoch was too late and we would have to increase the frequency of the digitally controlled VCXO in order to move the system epoch back into the correct position with respect to the reference epoch. In fact, if we made a fractional frequency increase exactly equal to the phase error measured in epochs, the phase error would return to zero one epoch after the frequency change was implemented.

If the Control Loop for the TKS were that simple, the VCXO output would be extremely noisy, because the phase meter has significant measurement errors each epoch. Instead, the measured phase errors are smoothed using a digital filter in order to determine the frequency correction to apply to the VCXO. We will now review the details of that process and point out the changes in the OB 8 software.

2.2 LARGE ERROR DETECT AND FAST/SLOW LOOP SWITCHING

Referring again to Figure 1, the computed phase error goes to the "Large Error Detect" function. If two consecutive phase errors greater than the "Failure Threshold" are detected, the software causes Non-Standard Codes (NSC) to be transmitted by the SV to prevent users from locking to the erroneous L-Band transmission. If this happens, it causes an outage of service, which can only be corrected by Control Segment intervention. The threshold for determining whether errors are large enough to cause NSC is the

"Failure Threshold," one of the fields in the SA Database Control Element (SADBCE). The NSC decision has two other effects. First, the Control Loop switches to a faster time constant (15 seconds), referred to as "Fast Loop." The original time constant will be restored when 50 consecutive errors less than a smaller threshold, the Clipping Threshold, have occurred. The mode, which uses the longer time constant, is referred to as "Slow Loop." The second effect of switching to the Fast Loop mode is that "Gain-Tracking" is disabled and ATC is disabled. The reason for this is that it is not possible to successfully adjust VCXO gain when large phase errors are present and the mean phase error is far removed from zero.

2.3 ADAPTIVE TIME CONSTANT (ATC)

If ATC is enabled via a Serial Magnitude Command (SMC) from the Control Segment (CS), the phase error is weighted by a function, which depends on the magnitude of the phase error. A number between 1 and 6 weights larger phase errors, and a number between 0.7 and 1 weights smaller phase errors (< 1 ns). Thus, the effect of ATC is to correct for large phase errors more quickly. Effectively, the time constant of the filter is shortened by a factor of about 6 if a phase error greater than a few ns occurs. In normal operation, the time constant selected by the Control Segment (usually 150-200 seconds) is used. Note that the ATC "weighting" does not occur if we are in "Fast Loop" mode.

2.4 LOOP FILTER

The phase error, possibly weighted by a factor between 0.7 and 6, is now applied to either the Fast Loop filter or the Slow Loop filter depending on the current mode of the TKS. The time constant in the Fast Loop mode is fixed at 15 seconds, while the time constant in the Slow Loop mode is derived from the SADBCE uploaded by the Control Segment. In OB 8.1 and following, there are more time constants available for selection by the Control Segment. In general, the ATC mode or shorter time constants are preferred if the VCXO is showing instability. The TKS starts with the mode set to Fast Loop, since the initialization process may result in initial phase errors, which need to be eliminated.

2.5 ADDING DESIRED DITHER AND SCALING BY RECIPROCAL OF VCXO GAIN

The output of the filter represents the desired frequency of the VCXO. However, to generate dither, a pseudo-random frequency change must be applied to the VCXO each epoch so the dither frequency is added to the filter output. After adding the dither, we have a number representing the desired fractional frequency offset of the VCXO from its center value, nominally 10.23 MHz. To get the VCXO to produce this desired frequency, however, we must know the frequency change produced by a 1 LSB change to the digital control input of the VCXO. If we know, for example, that the VCXO changes frequency by $2.5 \cdot 10^{-12}$ for each LSB change of the control word (the *VCXO gain*), then we must multiply the desired fractional frequency change by $4 \cdot 10^{11}$ ($1/[2.5 \cdot 10^{-12}]$) to compute the control word for the VCXO. This is shown by the multiplier function in Figure 1.

2.6 VCXO GAIN-TRACKING

To compensate for the fact that the sensitivity of the VCXO to changes in Control Word (i.e., the *VCXO gain*) may change over time, the TKS software continuously estimates the VCXO gain if certain conditions are met. First, the TKS must be generating dither, and second, it must be operating in Slow Loop. If the TKS is not generating dither, then it is not necessary for the VCXO slope to be estimated

with accuracy greater than about 5%, and the initial calibration of VCXO slope is accurate enough. This has been the case since 5/2/00 when dither was turned off.

3 COMPARISON OF TKS PERFORMANCE WITH AND WITHOUT ATC

To judge the effectiveness of ATC, the Orbital Buffer dumps are the most useful tool. This capability was first added with OB 8.0, which was uploaded on 10/10/00. These buffers hold 13 hours of 2-minute average phase error data, and must be captured via a commanded dump from the Control Segment. Since 10/10/00, these dumps have been captured at least once a week, and, once in a while daily or twice daily. Using these dumps, the TKS performance in the ATC mode has been compared to the TKS in the non-ATC mode with a filter time constants set at 90 s and at 150 s. The four time periods, which we will analyze, are:

1. **10/10/00 to 7/3/01** – TKS in 90-second mode. This period includes an eclipse and a large VCXO transient ($-7\cdot10^{-10}$) one day after eclipse onset, on 1/1/01. However, this event was only partially captured in the trace buffers and was missed in the 13-hour data. The effect of the large event is to cause a phase error of about 12 to 13 ns. We also see many, many smaller VCXO frequency steps, about half as large as the large "eclipse" events in this time period. These smaller frequency steps cause phase transients of up to –6 ns.
2. **7/3/01 to 11/28/01** – TKS in ATC mode. Towards the end of the eclipse (starting 6/29/01 and ending 7/27/01), and then continuing for another week, there are again many VCXO frequency steps of about $-3\cdot10^{-10}$ to $-4\cdot10^{-10}$. However, now the frequency steps only cause peak phase transients of about –2.2 ns.
3. **11/28/01 to 12/5/01** – TKS in 150-second (legacy, non-ATC) mode. We observe that the phase variation is reduced when ATC is enabled on 12/5. This is a stable time period for the VCXO, so we can see the normal phase variation caused by orbital temperature variation in both 150-second modes.
4. **12/5/01 to 3/10/02** – TKS in ATC (150-second) mode again. On 12/20/01 there is a "typical" eclipse-related VCXO frequency step ($-7\cdot10^{-10}$), the first in ATC mode, so we compare phase error in ATC mode to phase error in 90-second mode from the prior eclipse. There is a period of stable VCXO operation during this time period, so we can compare the normal orbital phase variation in ATC mode vs. 90-second and 150-second modes. We observe that well after the eclipse is over phase spikes caused by VCXO transients begin occurring, but seem to be receding at the end of the data. This allows us to compare TKS phase errors for 90-second vs. ATC mode, when the VCXO is having the smaller frequency transients.

3.1 TKS IN 90-SECOND MODE (10/31/00 TO 3/22/01)

Figure 2 shows the Orbital Buffer Graph for this period. First we will explain its contents, since there are many in this report.

3.1.1 UNDERSTANDING ORBITAL BUFFER GRAPHS (FIG. 2-19)

The blue (monochrome, dark) line represents the orbital buffer data unmodified. Each point is a 2-minute average of phase error with a resolution of 1/16 ns. Each buffer covers a 13-hour period. Over most of the period, the buffers were dumped on a weekly basis, so the lines between each buffer are of no consequence. At the bottom of the chart, in pink, is a single point representing a measurement of rms amplitude for the phase errors in *each* buffer. It is important to realize that:

1. Each point used to calculate the rms deviation of the buffer is a moving average of 15 points of 2-minute averages. The purpose in this is to reduce the "noise" caused by the phase meter and the noise introduced by the 1/16 ns quantization of the buffer data.
2. The rms values plotted at the bottom of Orbital Buffer Graphs will underestimate the true rms phase error of the TKS whenever spikes of less than 30 minutes' width.
3. The rms value plotted is based on a 12-hour period of half-hour-averaged data. Since most phase variations occur with a 12-hour period, this is more accurate than using all the data.
4. A value will not be calculated and a point will not be plotted if there is a partial buffer dump shorter than 12.5 hours.

3.1.2 Overview of the 90-second Data (10/31/00 to 3/22/01)

For the first 2 months of this 5-month period, the VCXO is well behaved, as is indicated by an rms phase error of less than 50 picoseconds ($5 \cdot 10^{-11}$). An expansion of a single orbital buffer on 12/6/00 is plotted in Figure 3. The phase error has a 12-hour cycle with only $\pm$ 50 ps of deviation. But beginning a week or two before eclipse, the average phase errors start increasing slightly, as can be seen by the larger deviations in the center of the chart as 12/31/00 is approached, and by the larger rms values in the bottom of the chart. The phase errors start increasing at about the time that eclipse season starts (dotted line in Figure 2) and increase gradually for about a month, then start to increase more rapidly. Finally, about 2 months after the beginning of eclipse, the errors reach a peak and then start to decay.

3.1.3 The Eclipse Anomaly on 1/1/01

However, one would have no clue that a major VCXO anomaly occurred on 1/1/01, or 1 day after the beginning of eclipse, by looking at this chart. There was no orbital buffer data dumped for that day, but the TKS almost went to non-standard codes (NSC) and the VCXO frequency apparently jumped by $-7 \cdot 10^{-10}$. This information was captured in the onboard TKS event buffers, which are designed to record phase errors in the event of abnormal conditions. In Figure 4 we see that the phase error probably reached an average value of about 13 ns, and a peak value above 14 ns. Two consecutive errors above 16 ns would have resulted in an "outage" – switching to NSC.

When this near-outage was first discovered, we thought the TKS was operating in 50-second mode. This led to the erroneous conclusion that we had just seen a VCXO frequency jump over two times larger than previous frequency jumps. That conclusion was erroneous, because the TKS was actually in the 90-second mode. Mode 3 of the TKS was changed from having a 50-second time constant to a 90-second time constant. OB8, uploaded on 10/10/00, replaced the original time constant table with new, more useful time constants, as well as adding ATC mode, the 13-hour orbital phase error buffer, and other TKS enhancements. There is actually no evidence that large eclipse events (frequency jumps) are getting any larger, since $-7 \cdot 10^{-10}$ jumps had been seen several times in the past.

3.1.4 The Post-Eclipse Phase Error Buildup

What is peculiar about the overall picture in Figure 2 is that the big VCXO frequency jump occurs when it is very well behaved, and only *after* the jump does the daily peak phase error start to gradually increase, indicating that the VCXO's frequency is becoming less stable. Figures 5, 6, and 7 show the orbital buffers 3.5, 4.5, and 5.5 weeks after eclipse. Rather than there being a wide phase error variation over the entire 12-hour period, the pattern instead consists of a positive frequency step (resulting in the positive phase triangle), followed 7.2 hours later by a negative frequency step (resulting in the negative phase

triangle). Although it cannot be seen in a single orbital buffer (13 hours), the asymmetric pattern does repeat (7.2, 5.8, 7.2, 5.8…) with a 12-hour period.

3.1.5 POST-ECLIPSE PEAK PHASE ERROR PERIOD

Figures 8 and 9 show the final buildup to the peak phase error and the peak phase error, respectively. The average phase error is 5.5 ns for the negative spike and 3.5 ns for the positive spike. It is clear that there is a negative and a positive spike caused by negative and positive frequency steps occurring every 12 hours for a period of about a month. The peak of the cycle occurs in this case somewhere between 53 and 65 days after the beginning of eclipse. There are apparently moderate-sized, negative frequency steps, which result in the negative phase spikes seen in Figures 8 and 9. Most significantly, these frequency steps *probably happen twice every day* for at least a 2-week period, since every phase buffer from 2/13/01 to 2/28/01 has them. The magnitude of the frequency step, based on the size of the spike (–6 ns, worst case) is about $-3.5 \cdot 10^{-10}$ fractional frequency change (dF/F). Similar-magnitude frequency steps can be inferred from later data in Figures 10, 12, 17, and 19.

3.1.6 SUMMARY OF TKS OPERATION IN 90-SECOND MODE

In this mode, the frequent, small frequency spikes cause about a 6 ns phase error (peak) and an rms phase error of about $4.5 \cdot 10^{-10}$. Large frequency steps (-7e-10) cause about a 13 ns phase error, and the largest VCXO events ($-10 \cdot 10^{-10}$) would probably result in an outage (NSC). To confirm this, examine Table 1, below, which summarizes the TKS response to all eclipse events.

When the VCXO is quiescent, the rms phase error, mostly due to orbital temperature variation is about $3 \cdot 10^{-11}$ (Figure 3).

3.2 FIRST TRANSITION TO ATC MODE ON 7/3/01

Figure 10 shows the average phase errors from 1 month before the transition to ATC mode to 2 months after the transition. The eclipse season had begun 5 days before the mode change on 9/29/01. We see a slow buildup of the peak phase errors occurring at 12-hour intervals. They finally reach –2.2 ns on 8/1/01, a week after the end of eclipse. It might appear from Figure 10 that ATC mode was responsible for this slow increase of phase errors, which peaked at the end of eclipse season. The phase errors peak in the range or –2.2 ns to +1.5 ns. However, comparing Figure 10 to Figure 2, we see a similar pattern, but with much larger phase errors. In these earlier data, the phase errors get as large as +3.5 ns to –6 ns. ATC mode, because it responds quickly to phase errors above 1.5 ns, actually reduced the peak phase errors by 63% for the peak negative error and 57% for the smaller peak positive error.

Figure 11 is a blowup of the phase error data covering just 5 days. It can be seen that the orbital phase variation, probably due to thermal effects on the VCXO frequency, is about 50% larger, as expected for the longer time constant. Figure 12 shows a detail of the phase spikes near the end of eclipse. Compared to the earlier phase spikes in 90-second mode detailed in Figure 9, they are similar in shape, but much smaller. Larger phase errors are suppressed by the nonlinear weighting of ATC.

3.3 TKS IN 150-SECOND MODE (11/28/01 – 12/05/01) AND TRANSITION TO ATC MODE

For a short time period when OB 9 was uploaded and before the SMC command, which enables ATC mode, was sent, the TKS operated back in the 150-second, legacy mode. This is the only time period in which the TKS was in 150-second mode for which we have orbital buffer data. We have already described how the TKS had outages in this mode whenever the VCXO had an eclipse anomaly, which happened with most eclipses except the first. Table 1 shows all the eclipse periods for SVN 43 and the effect on the TKS.

Figure 14 shows the transition from 150-second to ATC-150-second mode. It is displaying smoothed orbital buffer data (30-minute moving averages), and it can be clearly seen that ATC reduces the amplitude of the orbital 12-hour variation by about 30%. This effect is expected because larger phase errors are weighted more than smaller ones.

3.4 TKS IN ATC (150-SECOND) MODE (12/5/01 – 3/10/02)

There was a large VCXO frequency step on 12/20/01 of magnitude $-7 \cdot 10^{-10}$. Since ATC mode was in operation, the SV did not have an outage, but suppressed the effect of the frequency step. Figure 13 shows that after the step, the amplitude of the 12-hour phase variation increases, as it usually does during eclipse season. This is either due to a larger temperature variation, or greater sensitivity of the VCXO to the orbital variation. Figure 15 shows the 12/20/01 eclipse-event in greater detail. As with all eclipse events, the shape of the phase error indicates that the VCXO frequency change is *not* instantaneous, but builds up over about a 5-minute period.

The entire second time period during which the TKS operated in ATC mode is shown by the Orbital Buffer Graph of Figure 16. The pattern resembles the 90-second mode data in Figure 2, but with smaller phase spikes. Both Figure 2 and Figure 17 display phase errors that do not start increasing significantly until *after* the end of the eclipse period. On the other hand, the ATC-mode data of Figure 10 show similarly sized phase spikes, but the phase error starts building earlier, relative to the eclipse period. In Figure 16, the pink (light) bars on the bottom also show the gradual buildup of the magnitude of the phase spikes, but towards the end of the data, apparently not all 13-hour buffers have the spikes. A plot of the single largest phase spike is shown in Figure 17. When this is compared to Figure 15, we see that the moderate-sized, frequency steps are similar in shape to the major events, when you consider that ATC mode probably suppressed the "peakiness" of Figure 15. Finally, Figure 18 again shows that positive and negative phase spikes likely occur twice daily for periods of weeks, since every complete buffer with an rms measurement, has a similar, large value of $\sim 3 \cdot 10^{-10}$.

4 TKS PERFORMANCE COMPARISON FOR 150-SECOND, 90-SECOND, AND ATC MODES

We are now ready to compare performance of the three modes for which we have Orbital Buffer (13-hour) data. There are four criteria to judge the effectiveness of the TKS:

1. The TKS should never go to NSC when there is a temporary VCXO anomaly, if at all possible, and it should have sufficient margin so that (1) is true if the VCXO degrades;
2. The TKS should have a small orbital phase variation when the VCXO is stable, which is true at least 80% of the time;

3. The TKS should have a small peak phase variation when the VCXO is noisy and producing small frequency steps less than $4 \cdot 10^{-10}$, which apparently occurs four times a day for about 1 month out of 6.
4. The TKS should not have excessive short-term noise and, if possible, meet all Allan variance specs in the lab.

Table **2** summarizes all of these factors and how they are derived. The only alternative to consider would be ATC 200-second mode, but probably not for SVN 43. In retrospect the slight improvement in Allan deviation below 200 seconds is overruled by the better control of the ¼-day peak in Allan deviation that can occur because of the orbital temperature variation.

5 CONCLUSION

ATC mode has given reliable, low-noise performance. SVN 43 experienced five outages because of VCXO anomalies before ATC was enabled and none since. At the same time, the new orbital trace buffers have allowed us to determine that it now experiences lower phase variation due to orbital temperature variation than previous operation in the standard 150-second mode. Finally, very frequent phase spikes seem to happen four times a day for periods longer than a month, and the magnitude of these phase spikes, which can directly add to ERD errors has been reduced from 6 ns, or 2 meters, in 90-second mode, to 2 ns, or less than a meter in ATC mode. These spikes may well have been over 10 ns in the original 150-second mode. Similar improvements have been noted on SVN 46. We heartily recommend ATC mode to improve TKS performance for any SV with even the slightest VCXO instabilities.

Table 1. Eclipses, VCXO anomalies, and TKS response.

Eclipse Period	Event Date	TKS Mode	VCXO Event	Event rel. to eclipse start	Result
~ 2/6/98 – 3/4/98	none	150 s	none observed		No VCXO event – telemetry shows VCXO frequency was stable. A RAFS frequency jump occurred on 2/10/98.
8/2/98 – 8/31/98 (LOW BETA MODE)	8/16/98	150 s	Large frequency jump ($-10 \cdot 10^{-10}$)	+14 days	**OUTAGE** – double fast loop, NSC. A double fast loop is further evidence of a large frequency change in VCXO.
1/25/99 – 2/21/99	1/29/99	150 s	Large frequency jump ($-6 \cdot 10^{-10}$)	+4 days	**OUTAGE** – single fast loop (smaller anomaly), NSC.
7/21/99 – 8/20/99	7/22/99	150 s	Large frequency jump ($-10 \cdot 10^{-10}$)	+1 day	**OUTAGE** – double fast loop, NSC.
1/13/00 – 2/11/00	1/14/00	50 s	Large frequency jump ($-7 \cdot 10^{-10}$)	+1 day	TKS was unstable because of SW bug with gain tracking at 50 s loop, but outage avoided.
	4/6/00	150 s	Moderate freq. jump ($-4 \cdot 10^{-10}$)	+84 days	**OUTAGE** – single fast loop, NSC. May have been from spikes that occur after eclipse period (see Fig. 2,12) or may have been a consequence of Delta-V 1 hour before.
7/8/00 – 8/8/00	6/27/00	150 s at event date	Moderate freq. jump ($-5 \cdot 10^{-10}$)	–1 days	**OUTAGE** –single fast loop, NSC
12/31/00 – 1/28?/01	1/1/01	90 s	Large frequency jump ($-7 \cdot 10^{-10}$)	+1 day	Near-outage (15-20% margin). Fig. 4 shows –12.5 ns phase error. Surrounding orbital buffer dumps quiescent (Fig. 2).
	5/16/01	90 s	Moderate freq. jump ($-5 \cdot 10^{-10}$)	–43 days	Similar to event on 4/6/00 (between eclipses), but 90 s time constant prevented outage.
~6/24/01 –7/27/01	no "big" event	90 s, ATC after 7/3	Moderate freq. jumps ($-3.5 \cdot 10^{-10}$) probably twice daily	For 3 weeks before eclipse (90 s mode), then beginning again 2 weeks after it's start.	5 to 6 ns phase spikes visible in all of 3 13-hour orbital buffer dumps. Spikes become smaller (+1.5, –2.2 ns) after ATC mode.
12/19/01 – 1/19/01	12/20/01	ATC 150	Large frequency jump ($-7 \cdot 10^{-10}$)	+1 day	Orbital buffer (2 min. averages) shows –4.9 ns peak phase error. Trace buffer (Fig. 14) shows –5.7 ns peak phase error.

Table 2. Comparison of TKS performance in four modes.

Criteria	150-second	90-second	ATC (150)	50-second
Avoiding outages	poor – failed 5 times	perfect	perfect	perfect
Margin against outages with $-1 \cdot 10^{-9}$ frequency step[1]	none	none (–20%) (20% margin with $7 \cdot 10^{-10}$ freq. step)	50% (100% margin with $7 \cdot 10^{-10}$ freq. step)	25% (simulation)
Orbital phase variation (VCXO stable)	83 ps rms	30 ps rms	58 ps rms	not measured
Peak phase error with noisy VCXO (~$3 \cdot 10^{-10}$ frequency steps)	probably at least 12 ns	6 ns	2 ns	2.7 ns (simulation)
Short-term noise (Allan deviation)	good	fair	good – passed short term AV test in lab, but failed LT AV test. Fig. 14 shows superior on-orbit performance.	poor – noise from phase meter causes small freq. changes every epoch.

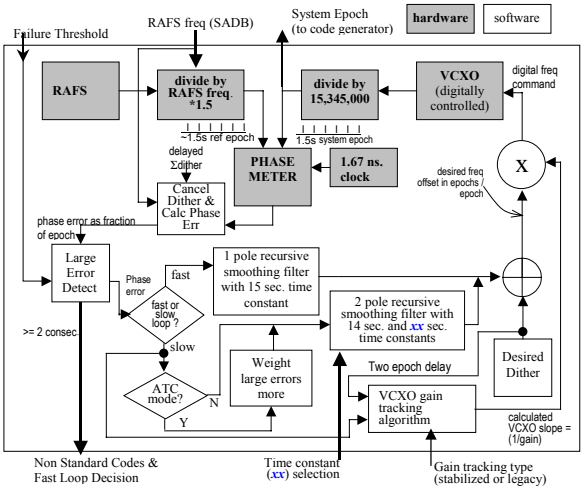

Figure 1. Block IIR TKS diagram.

[1] This is the largest frequency step that SVN 43's VCXO generated. It occurred twice. $-7 \cdot 10^{-10}$ is more typical. The step is not instantaneous.

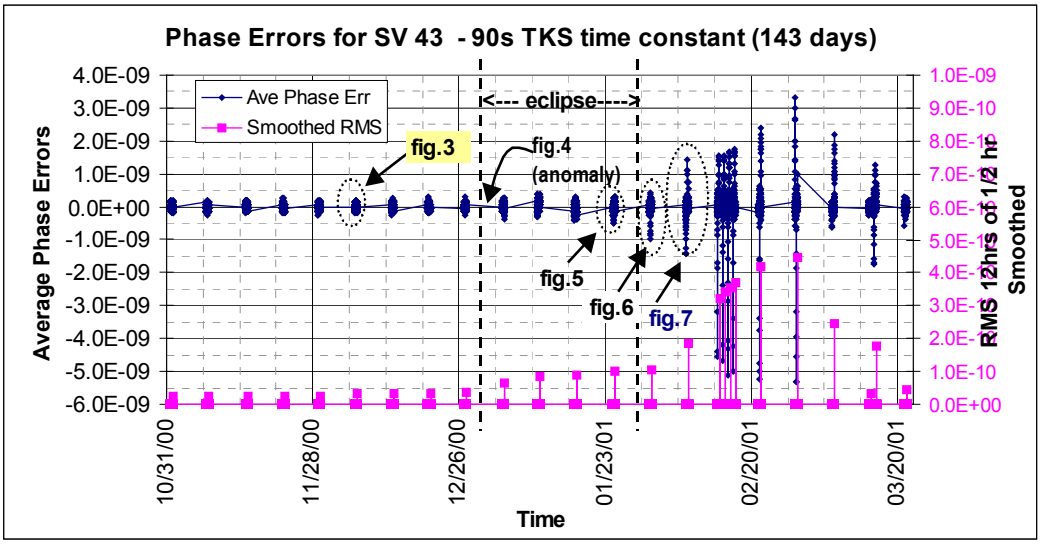

Figure 2. Long-term view of SVN 43 average phase errors (TKS in 90 s mode).

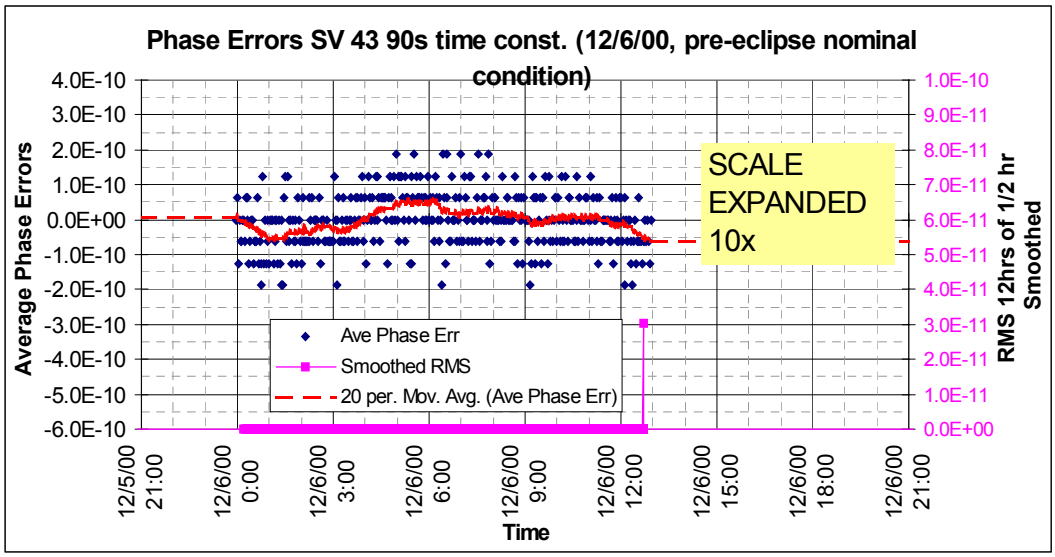

Figure 3. Typical data for quiescent VCXO condition with 90-second time constant.

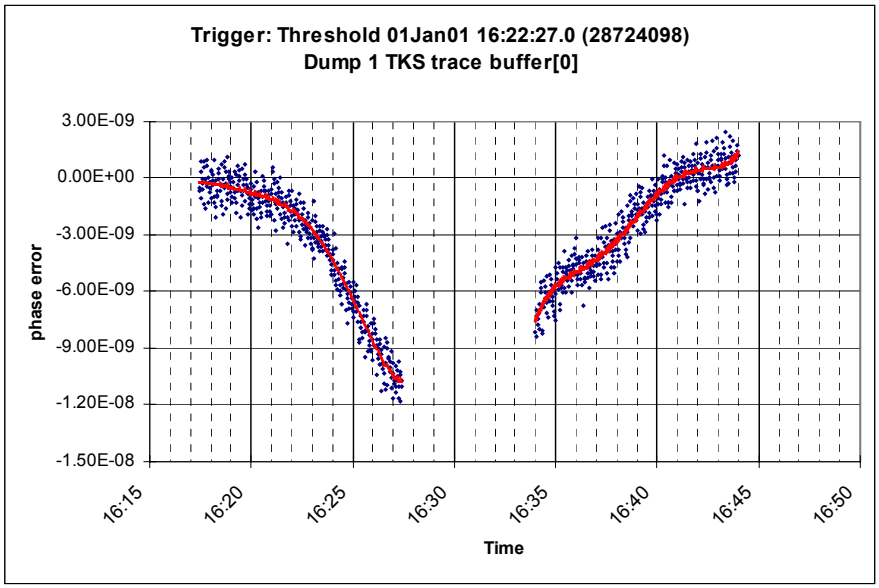

Figure 4. TKS Trace buffer data (1.5 sec) for VCXO $-7 \cdot 10^{-10}$ frequency step on 1/1/01.

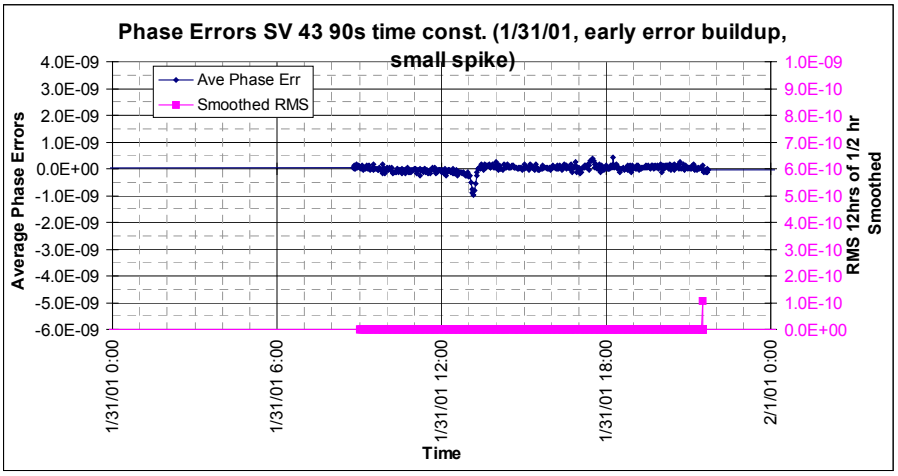

Figure 5. Buildup of TKS phase errors with 90-second mode after eclipse: 3.5 weeks.

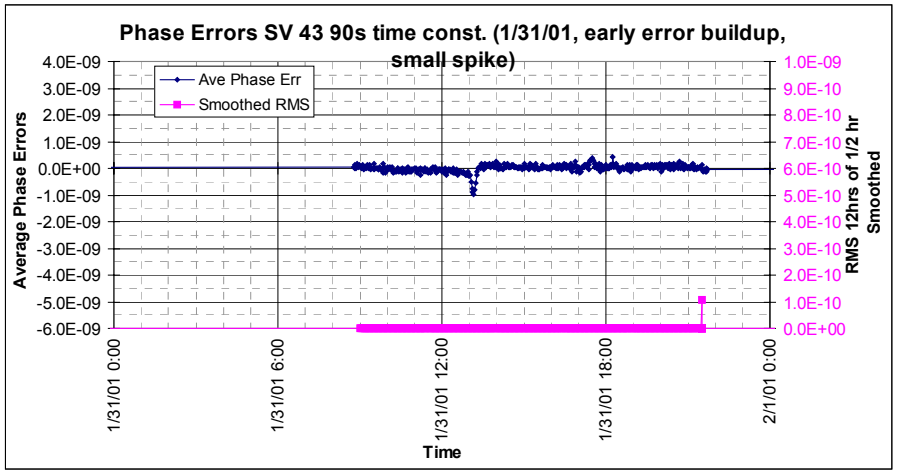

Figure 6. Buildup of TKS phase errors with 90-second mode after eclipse: 4.5 weeks.

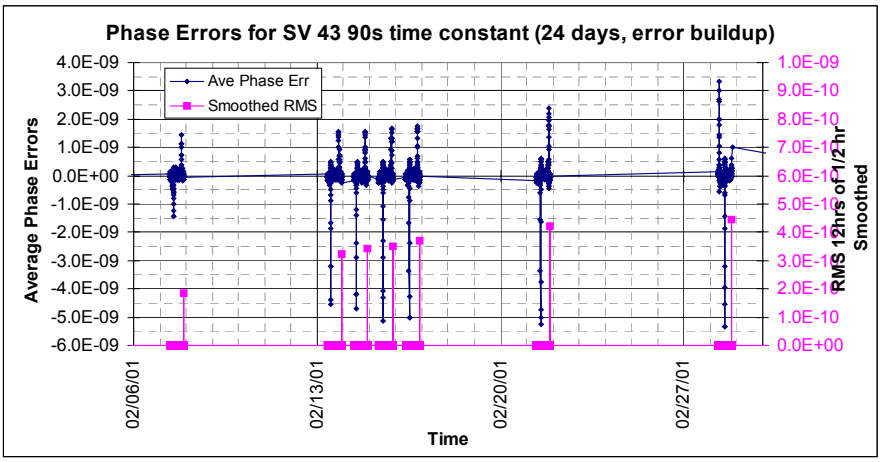

Figure 7. Buildup of TKS phase errors with 90-second mode after eclipse: 5.5 weeks.

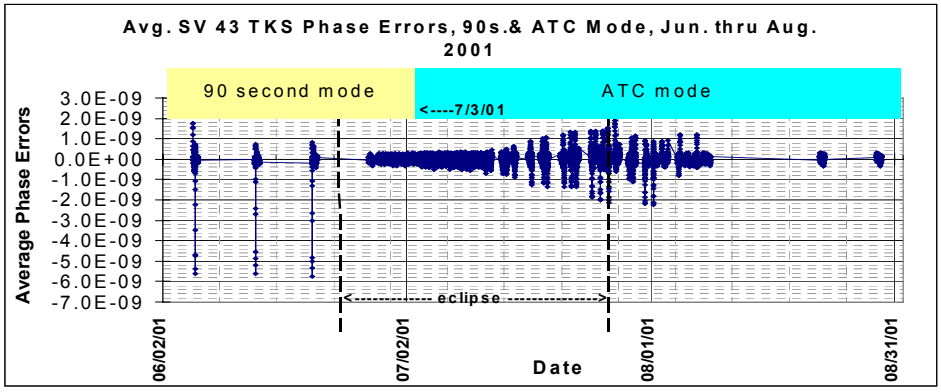

Figure 8. Final buildup of phase errors for 90-second TKS 1 to 4 weeks after eclipse end.

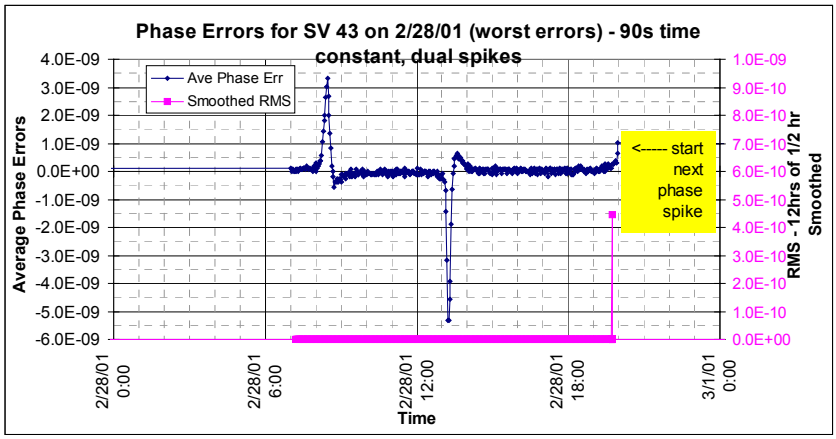

Figure 9. Detail of worst phase spikes on 2/28/01 for 90-second TKS 59 days after eclipse.

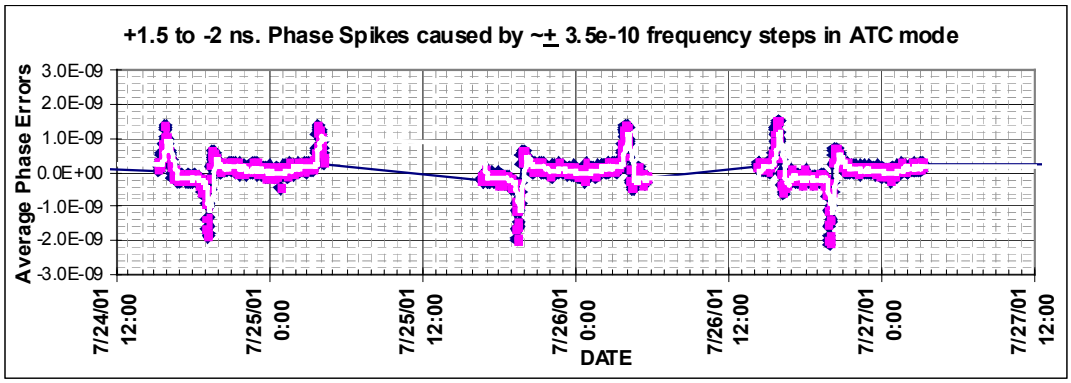

Figure 10. Phase spikes before and after transition from 90 s to ATC mode, 6/2/01-8/31/01. Moderate frequency steps cause 6 ns spikes in 90 s mode and 2.4 ns spikes in ATC mode.

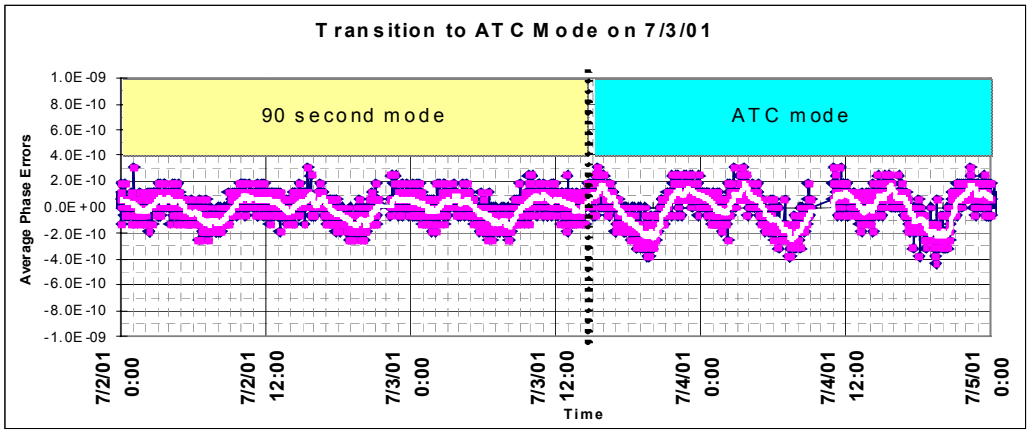

Figure 11. Detail of transition from 90 second to ATC mode on 7/3/01.

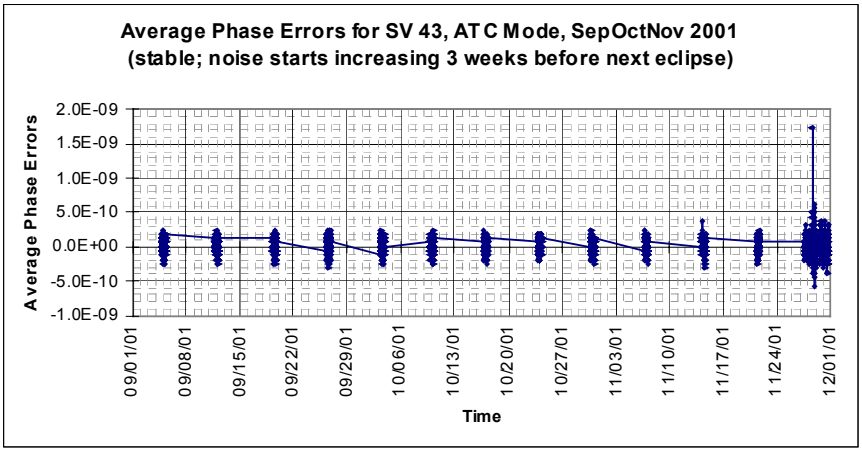

Figure 12. Detail of phase spikes at end of eclipse in Figure 10.

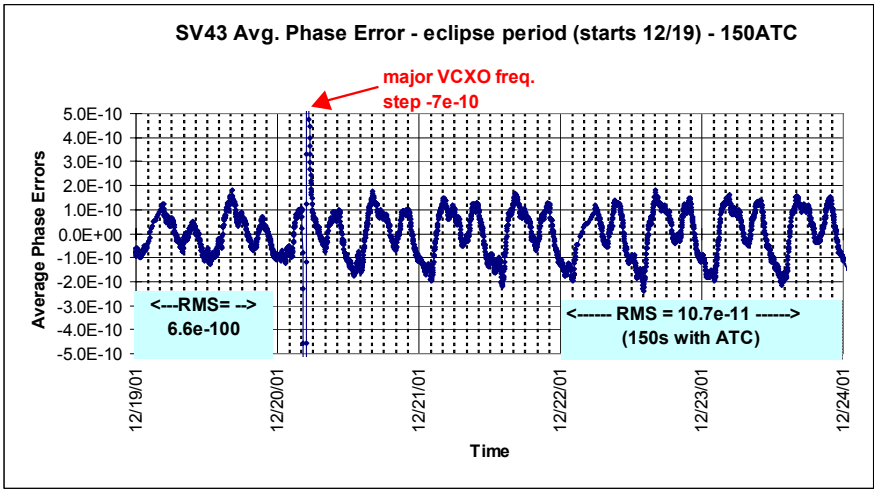

Figure 13. Next 3 months of ATC-mode phase errors. Spiking begins 3 weeks before next eclipse.

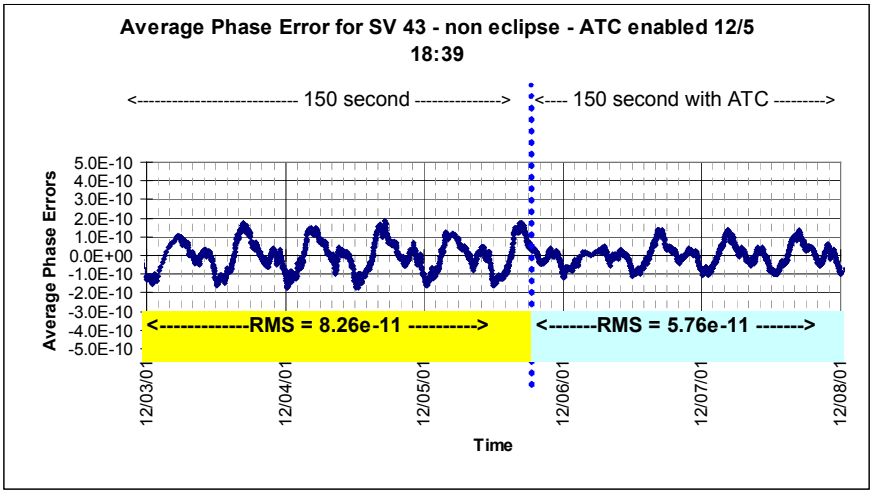

Figure 14. Effect of ATC on orbital phase variation (1/2 hour smoothed data).

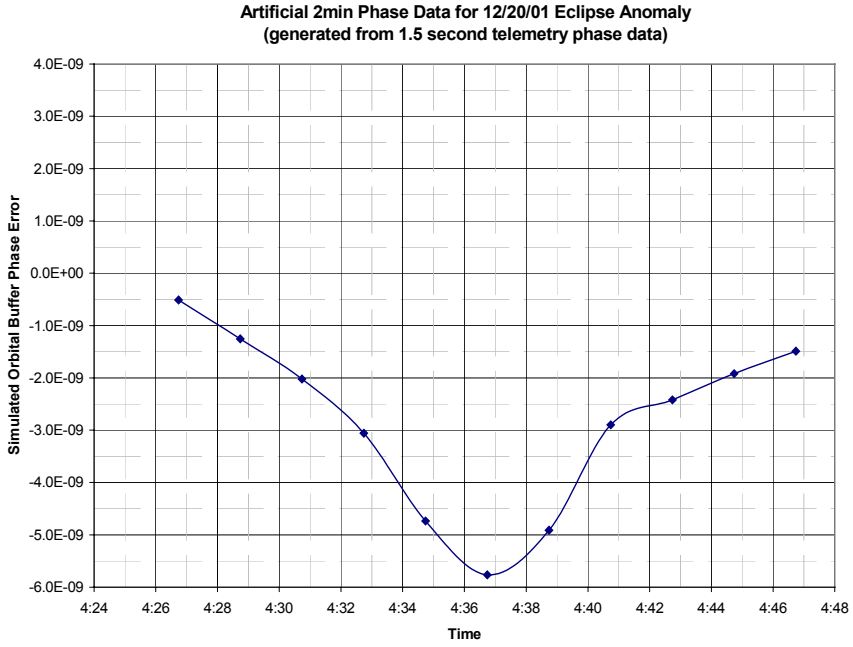

Figure 15. Increased phase variation after VCXO frequency step in ATC mode.

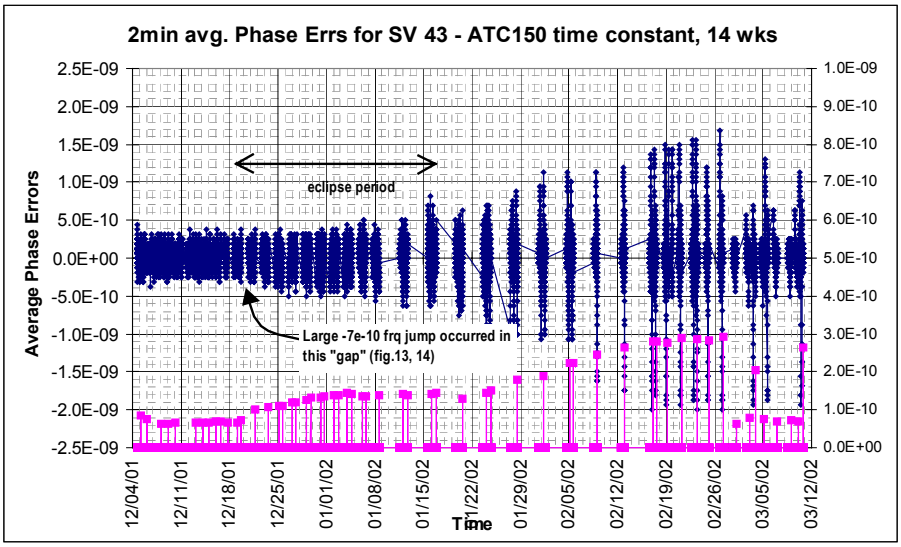

Figure 16. Orbital buffers (most of them) over second ATC mode period to 3/20/02.

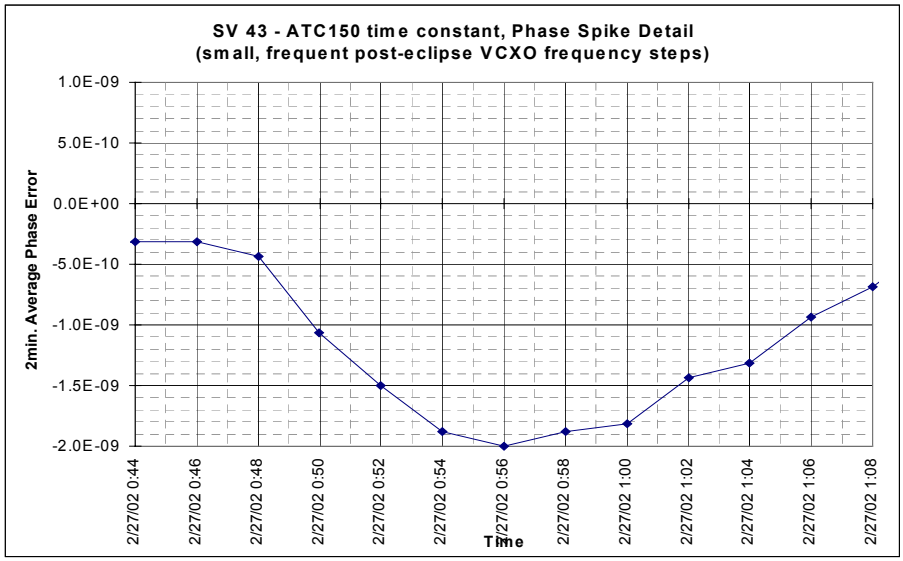

Figure 17. Largest ATC mode phase spike on 2/27/02 (–2.0 ns peak).

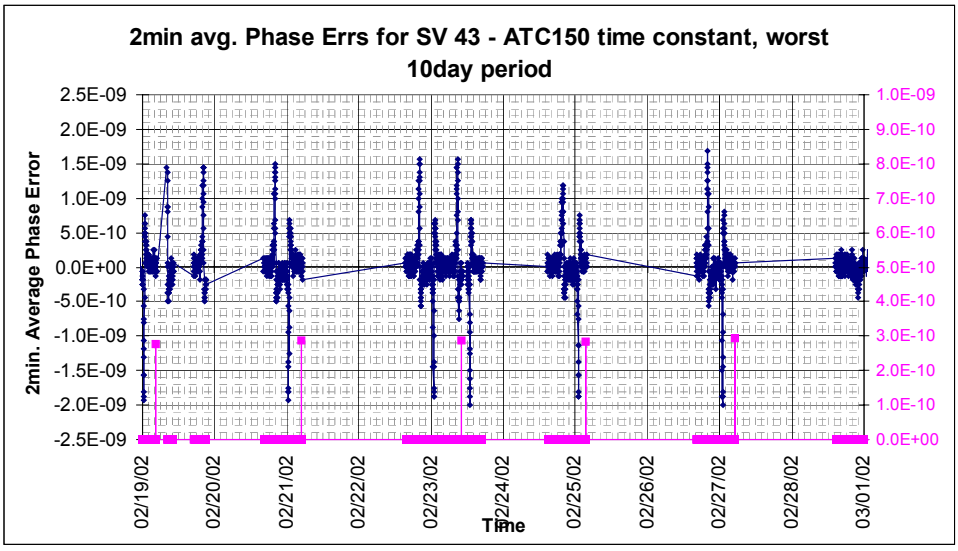

Figure 18. ATC mode phase spikes 2/19 to 2/27/02.

QUESTIONS AND ANSWERS

MARTIN BLOCH (Frequency Electronics, Incorporated): Did you simulate it on the ground? Or did the jump start in the orbit and you did not see it on the ground?

JOHN PETZINGER: We saw a few jumps on the ground, but it is kind of a thing that is hit or miss. We have some long-term tests; mostly we were testing the drifts, like for maybe a 20-day Allan variance Test, or something like that. And I kind of recall we did see a few where we got these unexplainable jumps, but they happen so seldom that we did not connect it with anything that was going to be a problem on orbit.

MICHAEL GARVEY (Symmetricom): The first slide you showed – if I read that correctly, is that sort of an open-loop quartz oscillator frequency?

PETZINGER: Yes, that is the open-loop – actually, if you look at that, that VCXO, that is an open-loop quartz oscillator frequency. And that is only varied by like $5 \cdot 10^{-9}$ in 4 years of operation.

GARVEY: I find it curious that the aging seems to be getting worse.

PETZINGER: Curious that it is aging so well? I mean, that is not drifting so much?

GARVEY: No, you show it zero for a while, and then it begins to become positive, I think. I did not look at it that carefully, but that is not standard performance.

PETZINGER: I ought to talk to you.

THOMAS PETSOPOULOS (U.S. Naval Research Laboratory): When you switch those time constants, doesn't switching the time constant or gain also cause a transient in your system? Or is that real fast and not noticed?

PETZINGER: No, because basically we are always near zero error, so changing the time constant doesn't induce an error; it just changes the rate at which you try to make the error zero.

JOHN PLUMB (University of Colorado): Can you explain why you think that is temperature- related? You say that happens during eclipses, but it seems to be pretty transient.

PETZINGER: Well, I think it is because of the fact that it happens at the space spikes that we see on the orbital variation when it is going through a certain temperature or something like that. We know crystals do these kinds of things, that they change state at some critical temperature.

PLUMB: And yet it seems that it is somehow magnified to the system. I am wondering, do you expect a temperature change of that magnitude to be able to induce a 6-ns error?

PETZINGER: Well, no. The worst case it does now is about 2.5 nanoseconds. And it is only for about a 10-minute period of time when it is doing that.