

MANAGEMENT OF PHASE AND FREQUENCY FOR GPS IIR SATELLITES

Dr. Marvin Epstein and Mr. Todd Dass
ITT Industries Aerospace, Communications Division
100 Kingsland Rd., Clifton, NJ 07014-1993, USA
Tel: 973-284-2614; Fax: 973-284-2351
E-mail: *Marvin.Epstein@ITT.com*

Abstract

The ITT-Industries-developed GPS IIR satellite payloads have been on orbit since 1997, providing outstanding signal-in-space performance. The GPS IIR satellites have the capability of adjusting the phase, frequency, and frequency drift of their broadcast clocks. This paper discusses the intricacies of phase and frequency management in the GPS IIR Time Keeping System.

During the initial years of the GPS IIR operation, the only clock correction technique used was the traditional technique of phase adjustment. In year 2000, the GPS IIR clock management operation started to use the full GPS IIR capability of adjusting phase, frequency, and frequency drift. Two approaches were employed. The first approach was applied to GPS satellites SVN43 and SVN46, where the phase, frequency, and frequency drift errors were zeroed during off-line maintenance. The phase and frequency of the other IIR satellites were managed by the second approach using on-line frequency drift adjustments to correct for phase, frequency, and frequency drift errors; and does not require the satellite to be taken out of service.

The operation of the two phase and frequency management approaches is compared in this paper. We have the results of over a year of operation of the new phase and frequency management techniques. As a result of these management approaches, the phase and frequency errors due to the GPS IIR rubidium frequency drift are now much smaller, with minimal load on the operation of the GPS system. This improvement is described in this paper by plots of the clock performance before and after the use of the phase and frequency management techniques. The cost of the improvement is minimal, as illustrated by a table of the ground control operations used to obtain these results.

INTRODUCTION

Each GPS SV clock must be synchronized fairly close to GPS time. Any residual offset between the SV clock and GPS time is broadcast to the user in the L-band Navigation message. The standard Navigation message format has a limited range of phase correction. In order to fall within the available phase values, the SV phase must be adjusted to be within ± 976 microseconds of GPS time. In addition, since the daily navigation upload for GPS IIR contains a 210-day prediction of the SV phase, the SV phase *forecast* is also required to stay within the ± 976 microsecond limit for 210 days following the upload, or the Navigation message cannot be built for the full span. After the SV phase is initially set, keeping it within these bounds is simple if the SV frequency and frequency drift residuals remain small.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE NOV 2001		2. REPORT TYPE		3. DATES COVERED 00-00-2001 to 00-00-2001	
4. TITLE AND SUBTITLE Management of Phase and Frequency for GPS IIR Satellites				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ITT Industries Aerospace, Communications Division, 100 Kingsland Rd, Clifton, NJ, 07014-1993				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001482. 33rd Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 27-29 Nov 2001, Long Beach, CA					
14. ABSTRACT see report					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

In all current generations of GPS SVs (Block II/IIA/IIR), phase and frequency are adjustable from the ground. A nominal GPS Block II/IIA cesium clock has a frequency drift that is on the order of 1×10^{-15} /day. Thus, a nominal cesium atomic standard, whose phase and frequency has been initially set to zero, would require zero or maybe one phase correction in its lifetime to keep the phase within the broadcast navigation message limits.

GPS Block II/IIA rubidium atomic standards have much higher frequency drift than cesium atomic standards. The GPS Block II/IIA rubidium atomic standards can have frequency drift on the order of 20×10^{-14} /day. Because of the high frequency drift, a nominal II/IIA rubidium atomic standard, whose phase and frequency has been initially set to zero, would require a correction to phase and/or frequency once or twice a year to keep the phase within the broadcast navigation message limits.

GPS Block IIR rubidium atomic standards, after a year of operation, have a frequency drift in the range of 1 to 5×10^{-14} /day. In addition, GPS Block IIR has an adjustable frequency drift that allows the Master Control Station to use a new frequency drift correction approach for Block IIR satellites to keep the phase within the broadcast navigation message limits. This approach has led to a much closer control of phase and frequency, so that frequency drift corrections are required, on the average, once every 1 or 2 years. There is a further benefit for the new approach in that a frequency drift correction can be applied to an operational satellite without any outage, while a satellite outage is needed for a phase and/or frequency correction.

During the initial years of the GPS IIR operation, the only clock correction technique used to control the clocks was phase adjustment. In year 2000, the GPS Master Control Station started to use the frequency drift correction approach.

The following sections will present our experience with the phase correction only technique, as well as the full reset (i.e. zero phase, frequency, and frequency drift) and the new frequency drift correction technique.

GPS IIR CLOCK ARCHITECTURE

A block diagram of the GPS Block IIR Time Keeping System (TKS) is shown in Figure 1. The source of the Block IIR signal timing for a given satellite is the on-board Rubidium Atomic Frequency Standard, which we refer to as the RAFS. The RAFS is a free-running clock at approximately 13.4 MHz, with no controls. The 13.4 MHz signal is passed through the on-board TKS, which controls a VCXO (Crystal Voltage Controlled Oscillator) to generate a 10.23 MHz transmitted clock, which is sent down to the ground to all the GPS users as an L-band signal. The 10.23 MHz signal can be adjusted in phase, frequency, and frequency drift by commands sent to it from ground control. GPS User Equipment can compute GPS Time from a given satellite by correcting the satellite's transmitted time by the clock residuals broadcast in the L-band Navigation Message.

COMMANDS

There are essentially three controls available to ground control to adjust the 10.23 MHz transmitted clock: 1) a phase adjust command, 2) a frequency adjust command, and 3) a frequency drift adjust command. The phase adjust command is quantized in cycles of the 10.23 MHz clock or approximately 97.8 nanoseconds, which is excellent for the task of centering the phase error close to zero. The frequency and frequency drift commands are quantized to 1.16×10^{-17} and to $1.02 \times 10^{-17}/\text{day}$, respectively. The quantizing of the controls for frequency and frequency drift are much finer than our ability to measure these quantities and, thus, meet our needs.

One strategy initially considered to keep the phase within the Navigation message limits was to use the three adjust commands to cancel out the clock phase, frequency, and frequency drift so that the phase would stay close to zero for the lifetime of the clock. This approach makes sense mathematically, but does not work out in practice because we are limited in our ability to measure frequency drift and because there are significant errors in our predictions of the future value of frequency drift. Once there are errors in frequency drift, the frequency drift will integrate to form frequency error and integrate again to form phase error. Because of these errors, one cannot just cancel everything and expect them to remain cancelled.

Instead, we decided to cancel most of the phase, frequency, and frequency drift to reduce the errors within some tolerance rather than do an exact cancellation. Given the tolerance and the management approach, one can see what is the cost and what is the benefit. A simple cost metric is the amount of effort needed to meet the tolerance (i.e. the number of uploads and outages required on the average). A simple benefit metric is the margin against some possible event that could cause us to overrun the phase error limits. Decreasing the tolerance by some amount will cause the cost and the benefit both to increase. We can compare different phase and frequency management approaches to select the preferred approach. We can also do a cost/benefit tradeoff to determine an acceptable tolerance.

GPS IIR EXPERIENCE WITH VARIOUS TECHNIQUES

In this section we will discuss our experience using three phase and frequency management techniques with the GPS IIR satellites. The three techniques are 1) phase correction only, 2) full reset (i.e. zero phase, frequency, and frequency drift), and 3) frequency drift corrections. Our experience is portrayed in Figure 2 and Figure 3, which are plots of the phase and frequency of all the GPS IIR operational satellites for the period from 1 January 2000 to 30 June 2001.

In the period from January 2000 to October 2000, the basic phase and frequency management technique was phase correction. During this period, the phase and frequency of SVN 43 (Space Vehicle Number 43) and SVN 46 covered a wide range of phases with peaks of over 800 microseconds and frequencies that covered the range from -31×10^{-12} to $+32 \times 10^{-12}$. Also, in this period the prediction of phase exceeded or was about to exceed the 976-microsecond Navigation message limit. The use of the phase correction-only technique led to the typical problems of high phase and high frequency error for rubidium clocks associated with this management technique.

In October 2000, the phase, frequency, and frequency drift of SVN 43 and SVN 46 were reset to zero. From the time of the reset to the present, the phase and frequency of SVN 43 and SVN 46 have remained very small, i.e. under 10 microseconds and 0.5×10^{-12} , respectively. Assuming nominal performance, *these two satellites will not need a phase or frequency management operation in the next several years.*

The other four Block IIR satellites, SVN 41, 44, 51, and 54, have used the frequency drift correction technique. These satellites were all initialized with a phase of about -100 microseconds and a positive frequency generally in the range of $+10 \times 10^{-12}$ to $+20 \times 10^{-12}$. SVN 51 was put into this initial state on 26 June 2000 after the satellite was operational. We see the transition to these values for SVN 51 in Figure 2 and Figure 3. The other satellites had this initialization before they became operational, and the initialization activity for these satellites does not show up in Figure 2 and Figure 3. The initial high frequency value decays to close to zero in 6 months or so and then the frequency stays small (i.e. between -5×10^{-12} and $+5 \times 10^{-12}$) for the rest of the plots. The phase for these satellites starts at -100 microseconds, decays some toward zero and then moves slowly, always staying inside the region of ± 200 microseconds. Up until now there has been only one frequency drift correction for each of these four satellites and it is anticipated that the average time between frequency drift corrections will be between 1 and 2 years.

Table 1 contains a history of the phase and frequency management activities of the Block IIR satellites. From the initial operation of the first IIR satellite in August 1997 until February 2000, no frequency drift adjustments were performed (i.e. the broadcast frequency drift was exactly the RAFS frequency drift) and the broadcast frequency was based on the original pre-launch estimate of the RAFS frequency. In February 2000, SVN 43 had an emergency frequency drift correction to address the failure of the phase forecast to be within the 976-microsecond Navigation message limits. In June 2000, SVN 51 had a phase and frequency adjustment after the SV was operational because we were developing our techniques and did not initialize the phase and frequency of SVN 51 before it became operational. Thus, we should discount the first three entries in Table 1 as part of the learning curve before we had developed proper phase and frequency management techniques. The next six entries correspond to the different techniques. In October 2000, we did a full reset on SVN 43 and SVN 46, with the SVs off-line. Then from 20 October 2000 to 24 July 2001, there were four frequency drift adjustments, namely one adjustment each for SVN 51, 41, 44, and 54. All the frequency drift adjustments were done without taking the SVN off-line. Thus, we can see that both the full reset and the frequency drift adjust approaches produce well-behaved phase and frequency curves with little effort. In the year 2001, we expect to have four frequency drift corrections for six satellites in orbit, which corresponds to a 1.5-year spacing between corrections. Similar results are anticipated in future years. Table 2 has the values of phase, frequency, and frequency drift on 30 September 2001 for the Block IIR satellites. It is expected that over the life of the clocks, we will get worse values than those shown in Table 2. However, Table 2 allows the reader to get a feel for the general level of values to be obtained with the proposed approach.

We looked at the two proposed approaches to see which approach was preferable. The full reset has a somewhat longer time between corrections. On the other hand, the frequency drift adjustment is much simpler to perform and can be done any time, as the frequency drift correction can be performed with the SV on-line. Our choice was to use the frequency drift approach, as we liked the ability to do a correction with an SV on-line any time and because the full reset correction might not last much longer than a frequency drift correction.

FREQUENCY DRIFT CURVE CHARACTERISTICS

The performance of our phase and frequency management technique for the Block IIR clocks depends on the frequency drift characteristic of the clocks and on our ability to predict the frequency drift behavior of individual clocks. This section will discuss what we generally know about the frequency drift behavior and the areas of uncertainty. The following sections will describe the management technique.

We have little knowledge of the future short-term frequency drift. Luckily, the frequency drift correction effectiveness depends on the long-term behavior of the frequency drift. By long term, we mean the average frequency drift over a month or number of months. Once the frequency drift is determined, then we can integrate it to get the frequency and phase.

The overall characteristics of the Block IIR RAFS frequency drift behavior are known from our experience with these clocks. For example, Figure 4 shows the early life frequency drift of SVN 54 over a 5-month period starting from about 10 days after clock turn-on. When the clock is first turned on, it has a very large negative frequency drift which moves towards zero. The magnitude of the frequency drift decays rapidly in the first few days after the clock has been turned on. When the SV is first tracked by the Control Segment Kalman filter, the frequency drift is in the range of -30 to $-10 \times 10^{-14}/\text{day}$. Within a few weeks, the average frequency drift is under $-10 \times 10^{-14}/\text{day}$. In the long term (e.g. after a year or more), our model of the frequency drift assumes that it is fairly stable and in the region of -5 to $-1 \times 10^{-14}/\text{day}$. Although the frequency drift decay rate is high at first, it rapidly gets smaller in the first few weeks after clock turn-on. Also, the frequency drift approaches zero, but always remains negative. This model matches the observations we have made of the Block IIR clocks.

Beyond the general model, we need specific predictions with specific numbers or ranges of numbers for individual clocks. Some clocks have frequency drift decay rates that are larger than others. Also, different clocks stabilize to different frequency drift values in the general range of -5 to $-1 \times 10^{-14}/\text{day}$, as described above. In addition, some clocks start their initial frequency drift decay with one curve and then switch to a slower curve. Because of these uncertainties, long-term predictions (i.e. more than a year) of frequency drift are not very accurate. Of course, we know what the frequency drift has been in the past and can make predictions that generally will be good for a year or so. Depending on the clock and the maturity of the clock, the predictions can be effective for much longer. In summary, we have some general ideas about the behavior of frequency drift, but there are limits to the accuracy of the predictions and the prediction error tends to grow with time.

FREQUENCY DRIFT CORRECTION PROCEDURE

This section describes our frequency drift correction clock management technique for a nominal clock.

This technique is split into three phases. The first phase is the start-up phase, where a clock has just been turned on and the SV is off-line, i.e. the SV is marked unhealthy for GPS user community. Next, there is the initial operation phase, during which the SV is marked healthy for the GPS user community. In the initial operation phase, the broadcast 10.23 MHz signal has the full frequency drift of the RAFS. The third phase is the normal operational phase, where the SV is on-line and the broadcast 10.23 MHz signal has a frequency drift which is significantly reduced from that of the RAFS.

In the start-up phase, a “cold injection” procedure is performed where the phase and frequency error of the clock are set to the vicinity of -100 microseconds and $+15 \times 10^{-12}$, respectively, and no reduction is made to frequency drift. In this phase, the broadcast 10.23 MHz signal has the full frequency drift of the RAFS. Once the cold injection procedure is performed, there need be no further corrections of phase or frequency in the life of a nominal clock. If there is some major disturbance to the clock, such as a change of the clock phase, then corrective measures may be employed including changes in the phase and frequency.

The initial operation phase starts when the cold injection procedure is completed and ends with the first frequency drift correction. In the initial operation phase, there are no corrections made to the clock, but there are changes in the phase, frequency, and frequency drift of the clock. The phase is driven from -100 microseconds toward zero, because the frequency is positive, which causes the phase to increase algebraically. The frequency is driven from $+15 \times 10^{-12}$ to zero, because the frequency drift is negative, which causes the frequency to decrease algebraically. The clock frequency drift starts off with a very large value such as -20×10^{-14} /day and decays rapidly toward zero and, at the same time, stabilizes. The frequency error becomes a cue as to when to make the first frequency drift correction to begin the next phase. When the frequency error is in the vicinity of zero, a frequency drift correction is made to cancel out most of the frequency drift of the RAFS. This leaves the broadcast 10.23 MHz signal with a negative net frequency drift, but much reduced in magnitude from the frequency drift of the RAFS. The initial operation period is on the order of 6 months long. At the end of the initial operation period, the frequency error is essentially zero, the phase is close to zero and the frequency drift is much reduced in magnitude from its initial large negative value and is much more predictable. The initial operation period is clearly shown in Figure 3, where clocks start off with large positive frequency and decay toward zero.

The normal operation phase starts off with the first frequency drift correction and continues for the rest of the life of the clock. During this time, frequency drift corrections are used to keep the clock frequency error between some limits such as $\pm 4 \times 10^{-12}$. Because the frequency drift of the clock varies, it is not possible to have the clock frequency drift correction exactly cancel the clock frequency drift. If the net frequency drift is positive, then the clock frequency error will go toward $+4 \times 10^{-12}$; if it is negative then the frequency error will go toward -4×10^{-12} . When the net frequency drift reaches a limit, a frequency drift adjustment is made so that the net frequency drift changes sign and the frequency error moves toward the other frequency limit. Typically, the frequency will have a zigzag curve between the frequency limits. *It is expected that the average time between frequency drift corrections will be between 1 and 2 years.*

The increase in phase error is fairly small because the frequency errors are kept to low values and because the zigzag frequency course tends to integrate into keeping phase low. The following heuristic argument illustrates this process. The positive frequency error values vary between 0 and $+4 \times 10^{-12}$, so that the average frequency error is close to $+2 \times 10^{-12}$. If the cancellation of the positive and negative frequency errors reduces the magnitude of the average frequency by a factor of two, then total average frequency error is about $+1 \times 10^{-12}$, which corresponds to an average phase error change of 31.6 microseconds per year, assuming 365.25 days in an average year. It will take many years (approximately 30 years) to go from 0-phase error to the maximum of 976 microseconds. If, for whatever reasons, we see a tendency for the phase error to get too large, we can adjust the frequency error limits to some other values (e.g. -2 to $+4 \times 10^{-12}$ or -4 to $+2 \times 10^{-12}$). These values will compensate for the average phase drift and move the average phase error toward zero.

SELECTION OF FREQUENCY DRIFT CORRECTION

A key element of this procedure is the selection of the frequency drift correction. To make this correction properly, one has to have a good estimate of the current RAFS frequency drift and the future frequency drift. Even estimating the current RAFS frequency drift is not simple, because there is a large amount of high frequency noise on the frequency drift. The tradeoff involved in the selection of a frequency drift correction can be illustrated by the following example. Assume that the frequency has reached the $+4 \times 10^{-12}$ limit and we have to select a frequency drift correction in the face of uncertainty of prediction of frequency drift. If the frequency drift correction is too small, then the RAFS frequency drift may wander so that the net frequency drift goes from negative to positive and we never reach the -4×10^{-12} limit. On the other hand, if the frequency drift correction is too large, then the net frequency drift will have a large magnitude and the -4×10^{-12} limit will be reached in a short time and we will have to make too many frequency drift corrections. The proper selection will reach the -4×10^{-12} limit in a fairly long time, over the expected range of variation of future frequency drift. Having predictable clocks with low frequency drift and slow variation makes it easier to manage the phase and frequency of the Block IIR clocks with only a few corrections over the clock lifetime.

The frequency drift procedure is basically robust in that it can provide the desired performance in the presence of perturbation and errors that are associated with operational conditions. It is expected that our frequency drift predictions will have errors, especially over a number of years, and that our decisions will be far from optimum when we examine them in hindsight. However, we have the benefit of making our new predictions based on what happened in the past and of learning from our mistakes. The result of these mistakes will typically be handled by a single extra frequency drift correction, which is a minor penalty. Given our expected margin over our goals, there should be little effort needed to meet our minimum expectation of one correction per satellite year.

CONCLUSIONS

We have a working phase and frequency management system that requires little effort and meets the requirements on phase error and forecast life. There is only a year or two of experience with this technique, so modifications may be made to take advantage of additional information as it becomes available. We would like feedback from the GPS community concerning any interest they may have with modifying the target ranges of phase, frequency, and frequency drift. While the reduction of the SV phase, frequency, and frequency drift residuals by this technique does not improve signal in space accuracy, it does improve the phase forecast and makes it easier to manage.

ACKNOWLEDGMENTS

The data for Figure 2, Figure 3, and Figure 4 were obtained from NIMA (National Imagery and Mapping Agency) at the Web site: <http://164.214.2.59/geospatial/products/GandG/sathtml/>.

The authors wish to thank the 2SOPS (2nd Space Operations Squadron) at Schriever Air Force Base for their support of the procedures described in this paper.




Figure 1. TKS Block Diagram




Figure 2. SV Phase from 1 January 2000 through 30 June 2001




Figure 3. SV Frequency from 1 January 2000 through 30 June 2001




Figure 4. RAFS Frequency Drift

Table 1. Phase and Frequency Management Activities

Date	SVN	Action	ON/OFF-line
8/97 → 2/00	All	Phase correction only	OFF
2/10/00	43	Frequency drift correction	ON
6/26/00	51	Set initial phase and frequency	OFF
10/5/00	46	Zero phase, frequency, frequency drift	OFF
10/10/00	43	Zero phase, frequency, frequency drift	OFF
10/20/00	51	Frequency drift correction	ON
3/1/01	41	Frequency drift correction	ON
3/1/01	44	Frequency drift correction	ON
7/24/01	54	Frequency drift correction	ON

Table 2. Phase, Frequency, and Frequency Drift for IIR Satellites on 30 September 2001

SVN	Phase (microseconds)	Frequency (10^{-12})	Broadcast Frequency Drift (10^{-14} /day)
41	-123	-0.22	+1.8
43	-6	-0.42	-0.2
44	-24	-3.21	-1.6
46	+4	+0.12	+0.2
51	-102	-2.95	-0.1
54	-75	-0.25	+0.2

QUESTIONS AND ANSWERS

AL KIRK (Jet Propulsion Laboratory): Just how exactly do you adjust the drift on these different clocks?

MARVIN EPSTEIN: There is a diagram that I skipped over, with the timekeeping system. But basically, we end up having a system that in effect says that we want the difference between the broadcast clock and the rubidium clock to match some curve. And there is a loop that works continuously inside the process to make sure it follows that curve. So now that curve has phase, frequency, and drift. So if I change the drift, in effect, I am forcing the broadcast clock to follow that drift.

KIRK: What mechanism do you use to adjust this drift?

EPSTEIN: The next thing we do to adjust the drift is that we upload a new drift value into a table, and then the software then calculates what changeover we should get given this drift. You have a curve going this way, now you want it to go that way. And then there is a closed loop that says, here's the error I see, here's our desired error, and it forces the VCXO, the broadcast clock, to go so it matches the desired curve. It's a closed-loop mechanism that is matching some software model of what's going on.

KIRK: I see. So, this is basically a software system that slowly adjusts the frequency physically on the clock.

EPSTEIN: It is a phase-locked loop that adjusts phase and frequency and drift. In other words, the model of the curve is a phase, frequency, and drift model, but the loop is basically a second-order phase-lock loop.

KIRK: Okay, thank you.

MICHAEL GARVEY (Datum TT&M): Is there a reason why you didn't use some of the more traditional clock techniques like Kalman filters?

EPSTEIN: Basically, what would drive the Kalman filter? What we are really doing is that we're looking at the ground, seeing what the clock is doing and then saying, get this drift and compensate it first for the next year and a half or 2 years, or whatever. And then the system just runs. There is no continuous feedback as to what you want or where you are going. All we're doing is just putting a number into a thing that says follow this curve. If we had a Kalman filter, what would it follow?

GARVEY: Well, it seems that you described a saw-tooth approach that was kind of: run to one limit and then head for the other, which, at face value, seems somewhat sub-optimal.

EPSTEIN: All we are trying to do is make sure we meet this 976 microsecond requirement over a long forecast time. We want to cancel out the drift. This mechanism seems to work, but if you have a Kalman filter, you are asking where would the input for the Kalman filter go. I find it is just managing the phase, what is wrong with the second-order loop, to track it.

When you build atomic standards, you have a Kalman filter track the output of the VCXO?

VICTOR REINHARDT (Boeing Satellite Systems): Actually, what I think you are doing is equivalent to an ARIMA filter integrated moving average kind of thing where you have a slow loop. So that works just as well as the other kind. In fact, I think you could show that they are equivalent.

EPSTEIN: What we have is a phase-lock second-order loop with a time constant that is only a couple minutes. And some human being is sitting there deciding what the frequency drift correction should be and puts it in there, and leaves it there for 2 years. So that is not a moving average, really.

DEMETRIOS MATSAKIS (U.S. Naval Observatory): Let me make a comment here. I think there is some confusion because I believe that the changes you make would not even be seen by users who rely on broadcast models. Whatever you do to the SV clock would be compensated by your broadcast corrections.

EPSTEIN: Right.

MATSAKIS: So people on the ground won't even know.

EPSTEIN: The correction is so smooth that before it builds up to a tenth of a nanosecond, we already have a new update in the satellite to go over and correct for it. The user never sees that.

MATSAKIS: So the only users who would see it are people who try to predict GPS clocks, particularly civilians.

EPSTEIN: That look at a raw clock and not the corrections.

MATSAKIS: I have actually done some of that for fun. You have asked for feedback. One thing that would be useful would be to be told well in advance whenever this is going to be done. So if there is a Web page, you could download past and future changes. That would be very good. Then do whatever you want, just tell us.

EPSTEIN: Okay, I think that sounds like a reasonable sort of thing. We will have to go to the Air Force and see what their rules are about making notification of these corrections. I think it is a very good thing to do. The question is that – NRL picks them up afterwards anyway. We do tell NRL about that when we do them. We do not really give a lot of advance notice. But you know, if we make a drift correction of two parts in 10^{14} per day, you cannot even see it from them – it is that small in that sort of time. You do not need advance notice of the thing. Because you try to figure out what you are doing, it's very hard to say.