Signal Deformations On Nominally Healthy GPS Satellites

Alexander M. Mitelman, R. Eric Phelts, Dennis M. Akos, Samuel P. Pullen, Per K. Enge, *Stanford University*

BIOGRAPHY

Alexander Mitelman is a Ph.D. candidate in the Department of Electrical Engineering at Stanford University. He received his S.B. in Electrical Engineering from the Massachusetts Institute of Technology in 1993 and his M.S. in Electrical Engineering from Stanford University in 1995. His research is focused on local area differential GPS design, signal analysis, and applications.

Dr. R. Eric Phelts is a Research Associate in the Department of Aeronautics and Astronautics at Stanford University. He received his B.S. in Mechanical Engineering from Georgia Institute of Technology in 1995, and his M.S. and Ph.D. in Mechanical Engineering from Stanford University in 1997 and 2001. His research involves multipath mitigation techniques and satellite signal anomalies.

Dr. Dennis M. Akos completed the Ph.D. degree in Electrical Engineering at Ohio University within the Avionics Engineering Center. He has since served as a faculty member with Luleå Technical University, Sweden, and then as Research Associate with the GPS Laboratory at Stanford University. Currently he is an Assistant Professor with the Aerospace Engineering Science Department at the University of Colorado at Boulder.

Dr. Sam Pullen received his Ph.D. in Aeronautics and Astronautics from Stanford University in 1996. Since graduating, Dr. Pullen has served as Research Associate and as Technical Manager at Stanford, where he has supported GPS Local Area Augmentation System (LAAS) and Wide Area Augmentation System (WAAS) research and development and is now the LAAS project leader. His work in these fields and his support of the Johns Hopkins University Applied Physics Laboratory (JHU/APL) GPS Risk Assessment earned him the ION Early Achievement Award in 1999.

Dr. Per Enge is a Professor of Aeronautics and Astronautics at Stanford University. He received his B.S. in Electrical Engineering from the University of Massachusetts at Amherst in 1975, and his M.S. and Ph.D., both in Electrical Engineering, from the University of Illinois at Urbana-Champaign in 1979 and 1983. Professor Enge's research focuses on the design of navigation systems which satisfy stringent requirements with respect to accuracy, integrity (truthfulness), time availability, and continuity.

ABSTRACT

This paper describes the collection and analysis of real data from twenty-four current, nominally healthy GPS satellites in the context of the ICAO second-order step (2OS) fault model. A simple but novel method is presented for averaging multiple C/A code epochs together, using Doppler prediction to achieve essentially open-loop synchronous sampling of the code waveform. This yields a clean picture of the C/A code edge transitions. Two different methods are then used to estimate the digital failure mode parameter, Δ, that characterizes each observed satellite within the context of the 2OS model. The results of this analysis are grouped by satellite block and, separately, by the type of frequency standard in operation at the time of measurement. Finally, the potential implications of these measurements on integrity and continuity are discussed.

INTRODUCTION

The Local Area Augmentation System (LAAS) and the Wide Area Augmentation System (WAAS) are specific applications of differential GPS to support aircraft terminal navigation and precision approach. Due to the safety-critical nature of LAAS and WAAS operations, the performance of these systems must meet stringent requirements on both integrity and continuity. In particular,

ground facilities must be able to reliably detect hazardous signal anomalies within a prescribed time-to-alarm while maintaining sufficient continuity for useful operation under nominal signal conditions.

One type of signal anomaly that is difficult to detect without specialized monitoring is the deformation of the C/A code signal broadcast by GPS satellites. The International Civil Aviation Organization (ICAO), through its GNSS Working Group B, has set forth a specific class of C/A code deformations against which any candidate DGPS-based precision approach and landing system must be able to protect its users [1]. This class of deformations is known as the second-order step (2OS) anomaly, and includes both analog and digital failure components.

To date, a number of theoretical analyses of the 2OS failure modes have been performed, and in our previous work we proposed a conservative upper bound for the analog failure component exhibited by a small subset of the operational GPS constellation [2]. On the other hand, few direct measurements have ever been made of the digital failure mode in healthy satellites. Because the edge distortion (jitter and/or delay) that comprises the digital anomaly on a healthy satellite is small by definition, making an accurate measurement requires tight control of the timebase in all recording equipment, as well as careful noise averaging over multiple epochs. The latter process is especially important when the RF chain has a larger-than-desired noise temperature, since amplitude variation due to thermal noise contributes directly to the estimate of jitter and delay on the code edges.

We motivate this paper with a simple calculation to illustrate the effects of small digital distortions on the two quantities of interest, integrity and continuity. Next, we briefly describe the experimental setup and introduce the methods used to estimate the digital mode failure parameter, Δ, from the raw data. Finally, we present the results grouped by satellite block and by frequency standard, and make some observations on trends evinced by the data.

MOTIVATION

One of the major challenges in designing a safety-critical system such as LAAS is the tension between integrity and continuity requirements. Integrity refers to the probability that the system will deliver hazardously misleading information (HMI) to a user without warning. For Category III LAAS, for example, the integrity specification is $P_{\mathrm{HMI}} \leq 10^{-9}/\mathrm{approach}$. Absent other considerations, the requisite level of integrity could be achieved by simply tightening alarm thresholds on the ground station. But the system also has a continuity specification, and this requirement limits the extent to which alarm thresholds can be tightened. Specifically, continuity is defined as the probability that an alarm will be issued to an aircraft on final approach. For Category III LAAS, the continuity specification for vertical guidance is $P_{\mathrm{Cont}} \leq 2 \times 10^{-6}/(15\,\mathrm{seconds})$. In this example, a small portion of that requirement, $P_{\mathrm{Cont}} \leq 10^{-7}/(\mathrm{independent\ test})$, is allocated to signal deformation monitoring.

To see why a small amount of delay on alternating C/A code edges can have a potentially significant effect on the landing system's ability to meet the stated performance requirements, we consider a signal with a $\Delta = 5$ nsec delay on the falling edges only, all other signal characteristics being ideal. Assuming a typical RF front-end (sixth-order Butterworth filter with 16 MHz bandwidth) on the ground reference station receiver, the differential position error is plotted as a function of Δ in Figure 1. The range error for $\Delta = 5$ nsec is approximately 0.75 cm. Given the Category III vertical alarm limit (VAL) of 5.3 m and conservatively assuming a poor satellite geometry (VDOP $= 5$), the error component due to the proposed edge delay is only VDOP $\cdot\, 0.75$ cm $= 3.75$ cm, or less than 1% of the vertical error budget. Thus a digital deformation of this size has negligible effect on integrity.

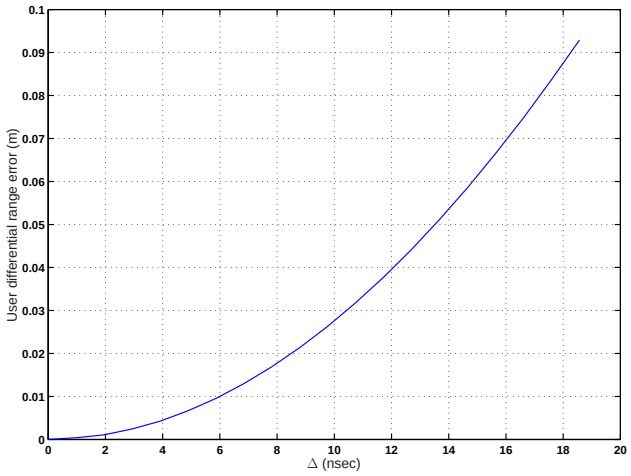

Figure 1: User differential range error as a function of Δ

On the other hand, the implications of a small falling edge delay for continuity may not be negligible. Within LAAS, the responsibility for issuing the alarms mentioned above typically rests with a module known as the Signal Quality Monitor (SQM). While a detailed description of the SQM algorithms used to generate alarms is outside the scope of this paper (see, *e.g.*, [1][3]), we consider one specific SQM function as a representative example. This function, or "metric," is simply a combination of several correlator

outputs from a reference station receiver:

$$D(\alpha_1, \alpha_2) = \frac{(I_{-\alpha_1} - I_{+\alpha_1})}{2 \cdot I_{\text{prompt}}} - \frac{(I_{-\alpha_2} - I_{+\alpha_2})}{2 \cdot I_{\text{prompt}}}$$

where $I_{-\alpha}$ and $I_{+\alpha}$ are the early and late outputs, respectively, of a correlator pair with a full-width spacing of 2α.

The SQM operates by continuously comparing an ensemble of such test metrics against precomputed "nominal" values, which are experimentally determined for each LAAS installation. Whenever one or more of the metrics exceed the nominal value by a predetermined threshold, the SQM provides this information to a module known as the Executive Monitor (EXM). The EXM, in turn, determines whether or not to issue an alarm based on the full set of information from SQM algorithms and other sources. The architectural and algorithmic details of a typical EXM are described in [4].

The statistical picture of our example test metric is shown in Figure 2. The threshold is computed for each test and each LAAS installation, such that the integral under the probability distribution curve and outside the thresholds meets the continuity specifications described above. This is customarily expressed in terms of the probability of false alarm: $P_{\text{FA}} \leq 10^{-7}$. For the example mentioned above,

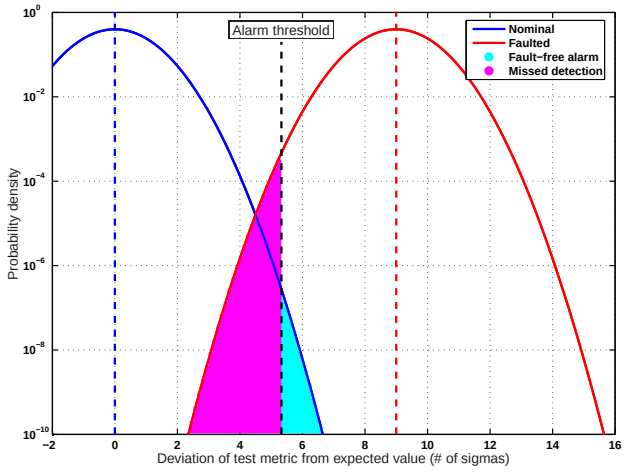

Figure 2: Probability calculations for a single SQM metric

and typical thresholds for the $D(0.1, 0.2)$ test, the presence of our proposed edge delay, $\Delta = 5$ nsec, increases the false alarm probability to approximately 1.22×10^{-7}, or a 22% increase. So even though a small edge distortion does not directly present any threats to integrity, it is nonetheless important to account for this effect if all LAAS requirements (continuity, in particular) are to be met simultaneously.

EXPERIMENTAL SETUP

In this section, we briefly describe the hardware and algorithms used for data collection.

Hardware Configuration

The hardware setup used in this experiment was substantially similar to that described in [2]. A high gain antenna was used to track each satellite under observation. The output of the RF chain from the antenna was fed into an Agilent 89600 vector signal analyzer (VSA), digitized, and stored to disk for post-processing analysis. The main changes were the use of a larger dish (46 m vs. 18 m diameter) and the addition of a high-performance external 10 MHz rubidium frequency reference in place of the internal oven-controlled crystal oscillator (OCXO) in the VSA (0.1 ppb vs. 200 ppb frequency accuracy). The rubidium source was calibrated against a primary cesium standard prior to the experiment.

Measurement and Processing Algorithms

One of the goals of this experiment was to explore the possibility of performing synchronous, high-resolution sampling and averaging of the C/A code waveform without having to do any demodulation or re-sampling in post-processing. To accomplish this, it was necessary to accurately predict the Doppler shift on a particular satellite at a specific GPS timepoint.

For each observed satellite, a current YUMA almanac [5] and the known location of the receiving antenna were used to calculate the expected Doppler shift at a particular instant of time several minutes into the future. Empirically, these predictions were surprisingly accurate: given a user position fix within 100 m (which is easily achievable with an inexpensive handheld GPS receiver), the predicted Doppler shift was consistently within 1.0 Hz of the measured value at the specified time point.

The predicted Doppler shift for each satellite was used to compute the effective center frequency and C/A code epoch period of the received signal, given by

$$f_{\text{c}} = f_{\text{L1}} + f_{\text{Dop}}$$

$$T_{\text{epoch}} = (1 - \frac{f_{\text{Dop}}}{f_{\text{L1}}}) \cdot 1 \text{ msec}$$

and these values were used to calculate the appropriate settings for the VSA. The instrument has a user-specifiable sampling rate (adjustable in steps of 1 mHz) with a maximum of $f_{\text{s,max}} = 47.5$ MHz, so the sampling frequency

used for a given satellite was set to

$$f_s = \frac{\lfloor T_{\text{epoch}} f_{s,\max} \rfloor - \frac{1}{D}}{T_{\text{epoch}}} \quad (1)$$

where D is a whole, positive number that represents the desired "Vernier ratio" for the sample set. This ratio is interpreted as follows.

Assuming the Doppler shift is constant throughout the recording interval, the sampling rate given in equation 1 yields an integral number of sample points in D code epochs. For our 250 msec-long data sets, this is a very good assumption: the change in Doppler over the duration of the set is typically less than 0.1 Hz.

The $D = 1$ case amounts to synchronous sampling; that is, the nth sample point in any particular C/A code epoch occurs at the same time within each epoch throughout the entire data set. This synchronization makes it possible to average multiple epochs together, reducing the effects of system noise (antenna noise and thermal noise from the RF chain) without blurring the underlying features of the C/A code waveform. This averaging yields the cleanest possible aggregate waveform from which to make Δ estimates.

Additionally, the $D > 1$ case can be used to achieve high-frequency sampling (that is, effective sampling frequencies above $f_{s,\max}$) while preserving synchronization. This approach, illustrated in Figure 3, directly trades time-domain resolution for noise averaging in a data set of a given size.

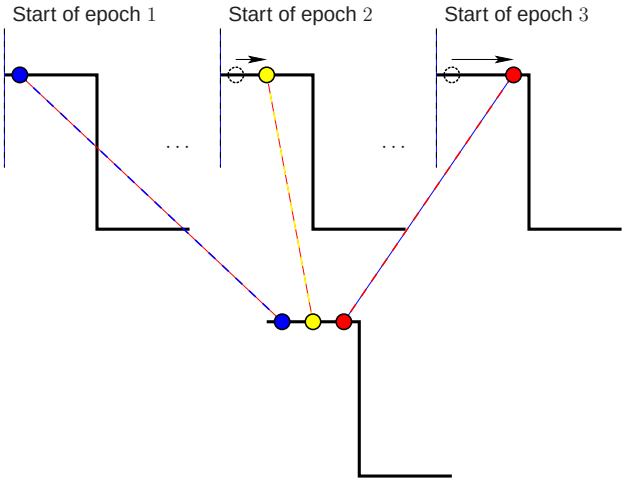

Figure 3: Vernier sampling by reducing the sample rate

As will be explained in the next section, it is also useful to sort the edges of each C/A code epoch according to the polarity of the extant navigation bit; thus the allowable values of D must take into account the relationship between code epochs and navigation bits. Since one navigation bit consists of exactly twenty code epochs, the useful values of D are the whole divisors of 20, that is, $D \in \{1, 2, 4, 5, 10, 20\}$.

For each observed satellite, multiple data sets $200 - 1000$ msec long were gathered using the procedure described above. After performing the process shown in Figure 3 and averaging the interlaced epochs together to reduce noise, estimates of edge delay and jitter were calculated from each resulting composite epoch using the methods described in the next section.

ANALYSIS METHODS

In this section, we describe two algorithms used to analyze the raw data gathered with the setup described above. The goal of each algorithm is to estimate the 2OS digital mode parameter, Δ, for each observed satellite.

Method 1: Direct estimation

The direct method of estimating Δ is based on the spacings between consecutive zero crossings in the C/A code waveform. Nominally, each spacing is an integral multiple of the C/A code chip period, corresponding to the number of consecutive positive (or negative) chips occurring at that point in the signal. Thus a plot of the spacings in an epoch of undistorted C/A code, modulo one chip period (corrected for Doppler shift, as described above), would look like Figure 4.

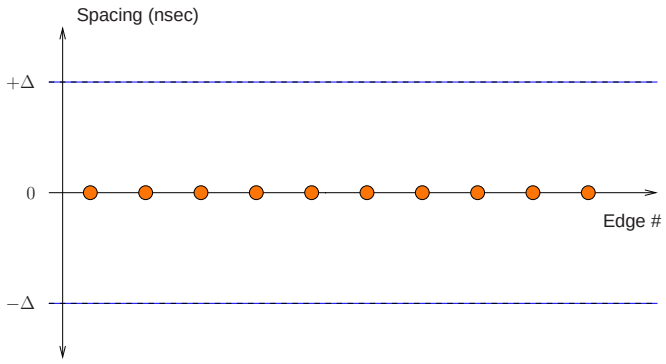

Figure 4: Spacing between zero crossings, modulo T_{chip}, under nominal conditions

On the other hand, a signal whose rising edges are nominal but whose falling edges are consistently delayed by some amount Δ with respect to their expected position would yield the result shown in Figure 5. While the definition of which edge polarity occurs at the "correct" (expected) time is an arbitrary choice, it does not affect the resulting plot. Furthermore, the sign of the spacing (indicating whether

a particular edge is early or late) is correctly determined by simply interpreting spacings (modulo T_{chip}) longer than $T_{\text{chip}}/2$ as being early relative to the next edge, rather than late relative to the current one. Thus a measured spacing of $3.9T_{\text{chip}}$ is interpreted as being $0.1T_{\text{chip}}$ early, rather than $0.9T_{\text{chip}}$ late. As long as the digital distortion is smaller than $T_{\text{chip}}/2$, this method will yield the correct result; distortion larger than that would be clearly visible by simple inspection of the correlation peak or other SQM methods.

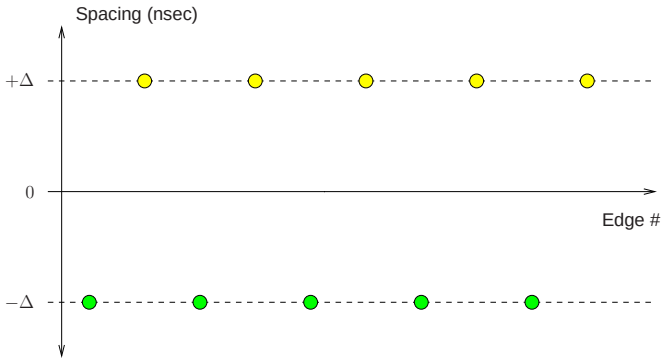

Figure 5: Spacing between zero crossings, modulo T_{chip}, with a delay of Δ on one polarity of edges (but not the other)

In estimating the edge delay as described above, we have ignored the effects of edge jitter as well as non-zero rise/fall times. Both will normally be present in any real system. With respect to jitter, this method will generate a usable estimate of edge delay provided that delay is significant relative to the amount of jitter. Additionally, a plot of the individual spacings between zero crossings will yield a direct picture of the jitter throughout the epoch. The use of zero crossings to estimate edge delay also implies that the results are completely insensitive to non-zero rise and fall times (whether due to a non-zero analog failure mode as described in [2] or for any other reason), *provided these times are symmetrical.* This is illustrated in Figure 6.

Finally, careful bookkeeping during post-processing makes it possible to determine whether digital distortion originates in the C/A code generator or somewhere downstream of it, for example in the navigation data unit (NDU). In order to make this determination, we sort each measurement by edge polarity (falling vs. rising) and by the sign of the navigation bit modulated onto each group of 20 code epochs.

Method 2: Indirect estimation

The indirect method of estimating Δ relies on the known properties of the C/A code pseudorandom sequence [6]. All sequences within the code family used by GPS satellites

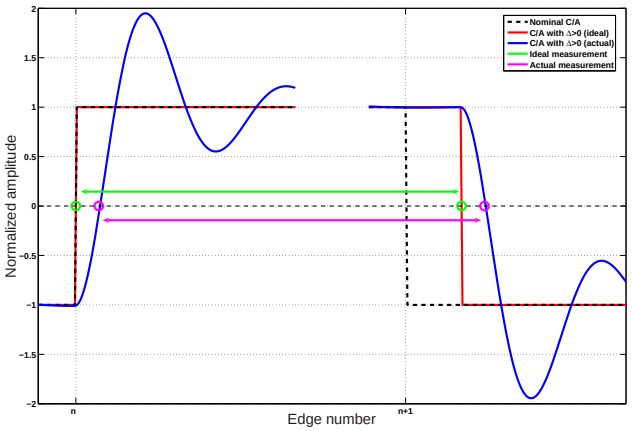

Figure 6: Symmetrical rise/fall delays have no effect on Δ estimate with the "zero crossing" method

are nearly balanced, so that each C/A code epoch contains exactly 511 chips of one polarity and 512 of the other. Because the effects of Doppler and analyzer sample rate are known to very good precision, it is possible to indirectly estimate edge delay by sorting all the samples in an epoch of data according to their sign.

Starting with a sample rate of f_{s} as defined in equation 1, the total number of sample points $c[\cdot]$ in a data set M epochs long is

$$N_{\text{total}} = MT_{\text{epoch}}f_{\text{s}}$$

Performing the post-processing step shown in Figure 3 on these samples yields M/D interlaced epochs of data with $N_{\text{chip, eff}}$ sample points per C/A code chip, where

$$N_{\text{chip, eff}} = \frac{D \cdot \lfloor T_{\text{epoch}}f_{\text{s, max}} \rfloor - 1}{1023}$$

Under nominal conditions, the balanced property of the code sequence implies that we should find exactly one chip's worth more samples of one polarity than the other in each epoch of data. In other words, if we define the quantity S as

$$S = \sum_{n=1}^{N_{\text{total}}} \operatorname{sgn} c[n]$$

then we expect $|S| = S_0 = MN_{\text{chip, eff}}/D$ for the whole data set when $\Delta = 0$. Any deviation from this indicates a non-zero Δ. We can use this fact to estimate the edge delay as

$$\hat{\Delta} = \pm\frac{(S/S_0 \cdot \operatorname{sgn} c[1]) - 1}{2(E_1 \text{ or } E_2)} \cdot \frac{T_{\text{epoch}}}{1023} \qquad (2)$$

where E_1 and E_2 are dependent on the particular PRN sequence being observed (specifically, they represent the number of sign changes in each epoch of code and whether

PRN	E_1	E_2
06, 09, 16, 28, 33, 34	256	256
07, 15, 17, 21, 24, 36	272	271
08, 22, 35	240	239
All other PRNs	256	255

Table 1: PRN-dependent coefficients for indirect method

or not the first and last chips in the sequence have the same sign). These values are given in Table 1.

Knowing whether to use E_1 or E_2 in equation 2 requires a priori knowledge of which edge polarity occurs at the "correct" (expected) time and which is faulted (delayed or advanced). In general, this is not possible for experimental data. Practically speaking, however, the effect of this ambiguity is small compared to measurement noise, so it does not significantly degrade the accuracy of the estimate. For the results presented in the next section, an average value $E_{\text{avg}} = (E_1 + E_2)/2$ was used in equation 2.

EXPERIMENTAL RESULTS

In this section, we present the results of our experiment and make a few observations about the data. These plots represent three data gathering sessions at the Stanford University 46-meter dish (February 2003, April 2003, and March 2004), using the hardware setup and algorithms described above; one session at Camp Parks, CA (September 2001) with no external frequency reference, predictive Doppler compensation, or synchronous sampling; and one session at Onsala, Sweden (January 2004), with a hydrogen maser external frequency reference and predictive Doppler compensation, but no synchronous sampling.

Individual satellite measurements

Examples of individual zero crossing measurements from two satellites, SVN 39 (PRN 09) and SVN 51 (PRN 20), are shown in Figures 7 and 8, respectively. The plot for SVN 39 shows no measurable edge delay, only a modest amount of jitter that would be expected from any real digital electronics module. Stated another way, all four combinations of navigation bit polarity and edge polarity yield measurements both above and below zero in the plot.

The plot for SVN 51, on the other hand, clearly indicates a persistent edge delay of approximately 3.75 nsec in addition to the expected jitter. Furthermore, the pattern of polarities contains an additional piece of information, namely that the fault appears to originate *after* the C/A code generator, presumably in the NDU. If the fault had occurred

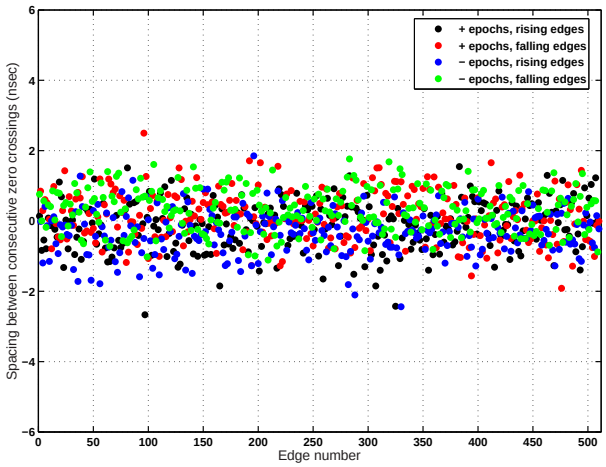

Figure 7: Zero crossing measurements for SVN 39

in the C/A generator, we would expect the falling edges that were delayed under one navigation bit polarity to have been advanced under the other bit polarity. Instead, this measurement shows the falling edges are consistently late, while the rising edges are consistently early.

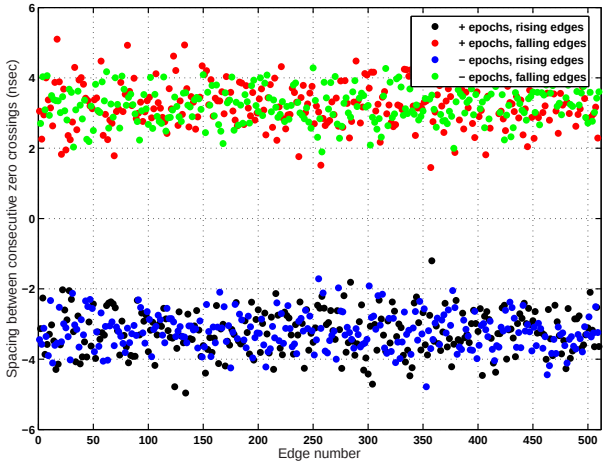

Figure 8: Zero crossing measurements for SVN 51

Collected measurements

Based on the methods described above, the average estimated values of Δ for all observed satellites are shown in Figure 9, grouped by GPS satellite block and plotted in ascending launch order. Within each block, there appears to be little or no trend with age. On the other hand, all of Block II-R exhibits $\Delta \geq 2.5$ nsec, with SVNs 41 and 47 (PRNs 14 and 22) in particular averaging an entire nanosecond larger than the rest. The outlying observation on SVN 41 is qualitatively consistent with the theoretical results presented in [7], although the magnitude of the distortion

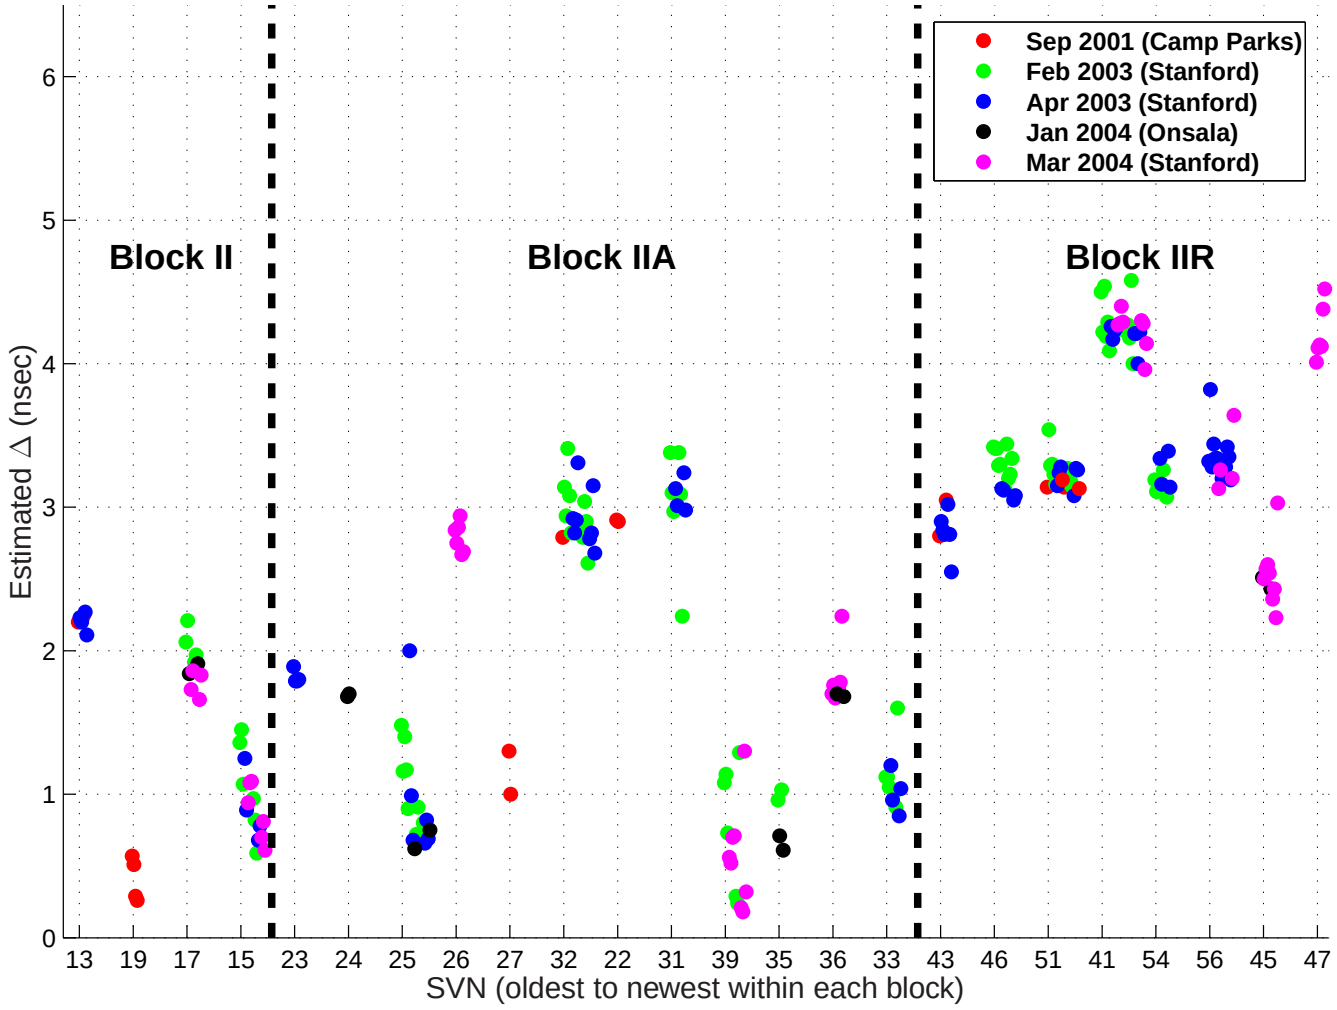

Figure 9: Delta estimates by satellite block

is a factor of five smaller in our measured data than in the earlier theoretical estimates. (SVN 47, launched in December, 2003, was not considered in the theoretical analysis.)

A slightly different organization of the data is shown in Figure 10, where the measurements have been grouped according to the frequency standard that was active at the time of observation. Here there seems to be a slightly more suggestive pattern. With the exception of SVNs 17 and 32, the satellites flying on rubidium frequency standards exhibit consistently larger Δ than those using cesiums. On the other hand, SVN 17 was measured both before and after its clock reference was switched from cesium to rubidium, and its estimated Δ actually improved slightly after the switch (from 2.1 nsec to 1.8 nsec). The suggested trend is therefore not absolute. As for SVN 32, one possible explanation is that its cesium clock was made by a different manufacturer (FEI) than the cesium standards in the rest of the satellites examined to date. Furthermore, this clock has been exhibiting small, periodic degradations in frequency accuracy over the period during which the measurements described in this experiment were taken [8][9]. To the extent that a satellite's Δ is affected by clock type, however, this case merits further investigation.

Finally, it is worth noting that the measurements presented above are consistent across three different observation sites, several variations of the experimental setup, two different numerical methods, and two and one-half years' time. Thus it is extremely unlikely that the observed phenomenon is in any way an artifact of the experimental procedure; rather, the digital distortion is real and almost certainly originates on the satellites themselves.

As demonstrated in the sample calculation, a non-zero Δ does not directly present a significant safety threat to users. On the other hand, it will be necessary to account for this effect in the error budget calculations for systems like WAAS and LAAS to ensure that they meet all stated performance requirements.

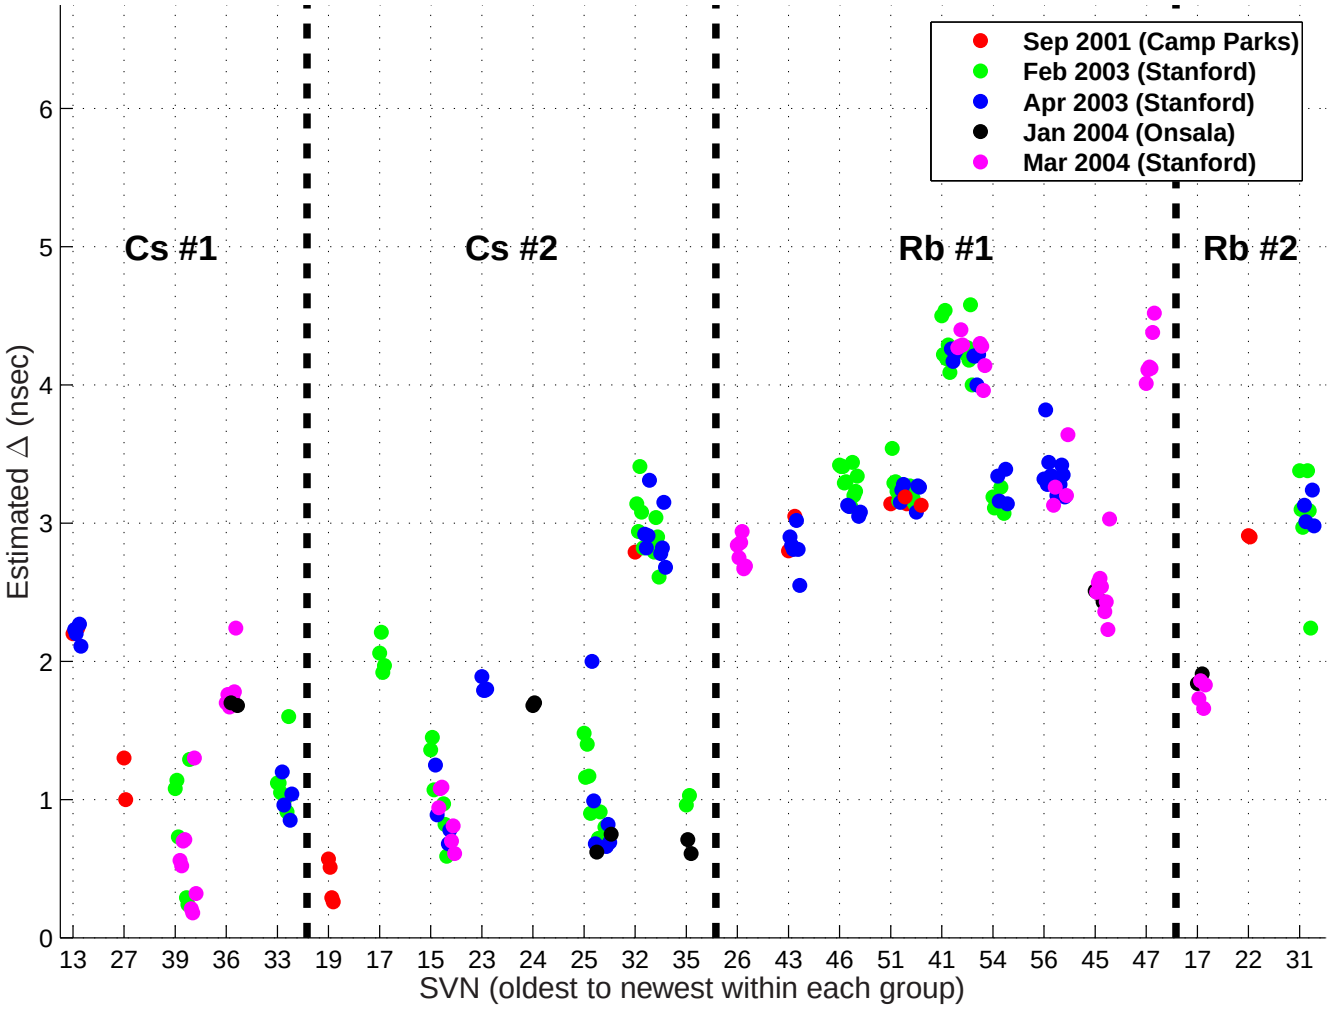

Figure 10: Delta estimates by clock type

CONCLUSION

The C/A code waveforms of twenty-four healthy GPS satellites were examined for digital distortion as defined by the ICAO second-order step threat model. A novel technique was introduced for synchronous high-resolution sampling and noise averaging of the code waveforms. Estimates of the digital distortion parameter were extracted from the raw data using two different numerical methods, and the results were grouped by satellite block and, separately, by frequency standard type. The results were found to be consistent across multiple facilities, experimental setups, and numerical methods over a period of two and one-half years.

Further study of this phenomenon should ideally include pre-launch testing of a GPS satellite, followed by periodic observations of that satellite throughout its operational lifetime. This would help determine the extent to which the digital distortion changes over time. Additionally, measurements should be taken of existing satellites after clock swaps, to further evaluate the effects of rubidium versus cesium frequency references on digital distortion. Finally, the numerical methods described in this paper may be extended to much longer averaging times, potentially enabling continuous monitoring of digital distortion without requiring a directional antenna. This information could, in turn, be used to generate and maintain a real-time database of failure mode probability allocations for safety-critical GPS-based systems such as LAAS.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Leonard Cutler and Mr. Jim Johnson at Agilent Technologies for their assistance, and the Federal Aviation Administration LAAS Program Office (AND-710) for their generous support of this research.

References

[1] Enge, P., R. Phelts and A. Mitelman, "Detecting Anomalous Signals from GPS Satellites," Working Paper 19, Global Navigation Satellite System Panel meeting, Toulouse, France, October 18-29, 1999.

[2] Mitelman, A., D. Akos, S. Pullen, and P. Enge, "Estimation of ICAO Threat Model Parameters For Operational GPS Satellites," *Proc. ION GPS-02,* Portland, Oregon, Sept. 2002.

[3] Shively, C., M. Brenner, and P. Kline, "Multiple Ground Tests Protecting Against Satellite Correlation Symmetry Faults in LAAS," Revision 3, RTCA SC-159, 1999.

[4] Xie, G., S. Pullen, M. Luo, P. Normark, D. Akos, J. Lee, P. Enge, and B. Pervan, "Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed," *Proc. ION AM-01,* Albuquerque, New Mexico, June 2001.

[5] United States Coast Guard Navigation Center, GPS Almanac Information Page, Internet URL: `http://www.navcen.uscg.gov/gps/almanacs.htm`

[6] Gold, R., "Optimal Binary Sequences for Spread Spectrum Multiplexing," *IEEE Transactions on Information Theory,* Oct. 1967, pp. 619-621.

[7] Brenner, M., P. Kline, and R. Reuter, "Performance of a Prototype Local Area Augmentation System (LAAS) Ground Installation," *Proc. ION GPS-02,* Portland, Oregon, Sept. 2002.

[8] Powers, E., U.S. Naval Observatory, personal correspondence, May 2003.

[9] Buisson, J., U.S. Naval Research Laboratory, personal correspondence, April 2004.