

The Role of the Clock in a GPS Receiver

Pratap N. Misra

Lincoln Laboratory,
Massachusetts Institute of Technology

satellites carry very precise, ultrastable atomic clocks costing tens of thousands of dollars each. The clocks used in most GPS receivers, however, are of an inexpensive kind for obvious practical reasons. Most applications do not ask much of a receiver clock beyond the ability to maintain a satellite signal in track, which translates into a requirement on phase noise in the clock signal. This requirement can generally be met with a temperature-compensated crystal oscillator costing in the neighborhood of \$100 or less.

The offset between the receiver clock and the GPS system clock ensemble is allowed to change with time without any significant constraints, though some receiver manufacturers control its size by adjusting or steering the receiver clock. The inexpensive clocks in receivers make GPS navigation practical, but, as is well known, introduce a minor complication: There is an additional unknown associated with each snapshot of measurements, namely, instantaneous receiver clock offset or bias relative to GPS (System) Time, the estimation of which requires an additional satellite measurement. That is why four satellites are normally needed for three-dimensional positioning rather than three.

In this article, we try to answer the question, What benefits would accrue to a GPS user with a "perfect" receiver clock? But before we can answer this question, we must first decide what we mean by *perfect*. Obviously a clock keeping time in synchronism with GPS Time would be perfect. But that's actually a clock with more accuracy than we need. Consider a clock with a known or precisely predictable offset at any instant relative to GPS Time. It's easy to see that such a clock would also be perfect — perfectly predictable. The main requirement is that the instantaneous clock bias be known somehow. The problem, as we'll see, is not simply that the estimation of the clock bias "uses up" a measurement, but that it alters the basic

structure of the position estimation problem in a negative way.

A user with a perfect clock would obviously be able to continue navigating with GPS even when the number of satellites in view falls to three. But that's the least of it. A perfect receiver clock can offer significantly better vertical position estimates, a built-in test of the integrity of the satellite constellation and the position estimates, and more robust use of carrier-phase measurements for precise positioning with easy detection of cycle slips and quicker resolution of integer ambiguities. And, it turns out, the benefits derived from a perfect receiver clock can be achieved substantially if a clock meets the simple requirement that its bias relative to GPS Time changes "smoothly." The idea is straightforward: A smooth function is predictable over a short interval, given the recent trend.

In the next section, we will look in detail at GPS positioning with a perfect receiver clock. We'll follow that with an examination of the correlations that are observed among the position and clock bias estimates obtained from a snapshot of GPS measurements. The results of this analysis are used to define an approach based on clock modeling to improve the quality of vertical position estimates. Finally, we will look at the benefits of a predictable clock in integrity monitoring and in carrier-phase processing.

HOW PERFECT IS PERFECT?

Given a snapshot of n GPS pseudorange measurements ($n \geq 4$), we can estimate the four unknowns: the three user position coordinates (x, y, z) and the instantaneous receiver clock bias (b). We'll refer to this approach as *4-D estimation*. On the other hand, if the receiver clock bias relative to GPS Time were known precisely, GPS-based navigation could proceed by estimating only the position coordinates. This will be referred to as *3-D estimation*. We know *a priori* that having one less parameter to estimate from the measurements should improve the quality of the remaining estimates. The question is by how much?

We examine first the quality of the position estimates with and without a perfect clock (3-D versus 4-D estimation) at a root-mean-square (rms) level. The idea is to base the comparison on the following relationship:

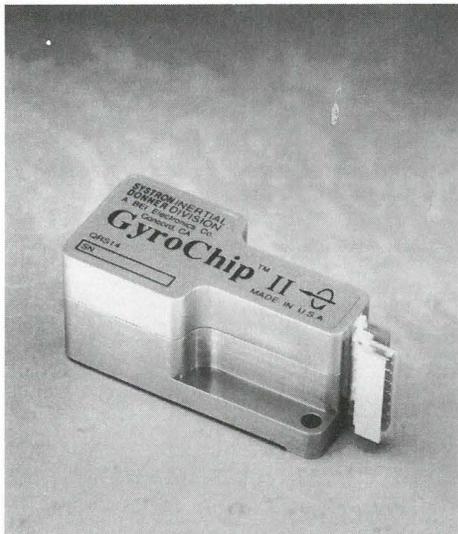
$$\text{rms position error} = \sigma_{\text{URE}} \times \text{DOP}$$

where σ_{URE} is the rms value of the error in the user's pseudorange measurements, and *DOP* denotes *dilution of precision*, a param-

It sounds a little strange but the most precise way of measuring a distance is to use a clock. Time, the quantity most difficult to define, is the one we know how to measure most precisely. In fact, the length of the meter is defined in terms of the length of the second through the adopted value of the speed of light in a vacuum. It is this fundamental relationship (distance = speed \times time) that is at the heart of how GPS works. By measuring the time elapsed for a signal to propagate from a satellite to a receiver and multiplying it by the speed of light, a GPS receiver can determine the range to the satellite. But there's a hitch. Any error in the time-keeping capability of the receiver's clock will be reflected in the computed range. In this month's column, Dr. Pratap Misra, senior staff member of the Massachusetts Institute of Technology's Lincoln Laboratory in Lexington, Massachusetts, will review the role of the clock in a GPS receiver and the effect its performance has on GPS position accuracy.

"Innovation" is a regular column in GPS World featuring discussions on recent advances in GPS technology and its applications as well as on the fundamentals of GPS positioning. The column is coordinated by Richard Langley and Alfred Kleusberg of the Department of Geodesy and Geomatics Engineering at the University of New Brunswick, who appreciate receiving your comments as well as suggestions of topics for future columns. To contact them, see page 4 of this issue, under the "Columnists" section.

As every newcomer to GPS learns, clocks are at the heart of this technology. The GPS



GyroChip II™

Solid-State Rotation Sensor

Get the performance and reliability you need for measuring angular velocity, in a compact, light-weight, low-cost package. The GyroChip II is built to last in demanding environments, with built-in power regulation and DC-in, DC-out operation. You can also select the GyroChip II model that's right for your application: standard, for battery systems and single-sided (+12 V) power supplies, and low-noise, for double-sided (± 15 V) supplies.

Put the GyroChip II to work:

- Cover GPS interrupts
- Improve survey resolution
- Add real-time navigation capability
- Improve attitude determination

Only the GyroChip II comes with our decades of experience in instrument design and manufacture.

If It Moves, GyroChip It!



SYSTRON DONNER INERTIAL DIVISION
BEI Sensors & Systems Company

2700 Systron Drive
Concord, CA 94518
USA: (800) 227-1625
Customer Service: (510) 671-6464
FAX: (510) 671-6647

European Business Office
Phone: 44 1303 812778
FAX: 44 1303 812708

Circle 39

ter characterizing the satellite geometry. Because σ_{URE} is the same in both 4-D and 3-D estimations, we need only compare the distributions of the DOP parameters in the two cases. The definitions of HDOP and VDOP (horizontal and vertical DOPs) in terms of the observation matrix ($4 \times n$) for the linearized problem of 4-D estimation are well known (see "The Mathematics of GPS" in the July/August 1991 issue of *GPS World* for a discussion of various DOPs and the linearization problem). For 3-D estimation, these parameters are redefined appropriately on the basis of the revised observation matrix ($3 \times n$) so that the relationship in the equation above remains valid. We will refer to the DOPs for 3-D estimation as HDOP3 and VDOP3 to distinguish them from the established notation in the 4-D case.

Figure 1 gives the cumulative (probability) distribution functions (CDFs) of the HDOP, VDOP, HDOP3, and VDOP3 available to users globally from the 24-satellite GPS constellation with a 5-degree elevation angle mask. The VDOPs indeed tend to be larger than the HDOPs, confirming the widely held view that vertical error in GPS-based position estimates tends to be larger than horizontal error. But the CDFs of HDOP3 and VDOP3 tell a different story for 3-D estimation: VDOP3 shows little variability, and its values are significantly smaller than those of VDOP and somewhat smaller than those of HDOP; and the distributions of HDOP3 and HDOP are substantially similar. In other words, a GPS user equipped with a perfect receiver clock would obtain much-improved vertical position estimates through 3-D estimation because of the much-improved DOPs. By contrast, there would be little gain in horizontal accuracy.

Clearly, the price for estimation of the clock bias is being paid for almost entirely by loss of accuracy in the vertical dimension. But why? And how much of this loss can be recovered with a practical receiver clock?

CORRELATIONS OF 4-D ESTIMATES

We examine next the relationships between GPS-based estimates of position (x, y, z) and clock bias (b) obtained through 4-D estimation. Figure 2 shows scatter plots of errors in the horizontal and vertical position estimates versus the error in the corresponding clock bias estimates computed from GPS pseudorange measurements, snapshot by snapshot, over a day at a surveyed antenna location using an atomic cesium beam oscillator for the receiver clock.

Figure 2a is a plot of the errors in estimates of a horizontal position versus the

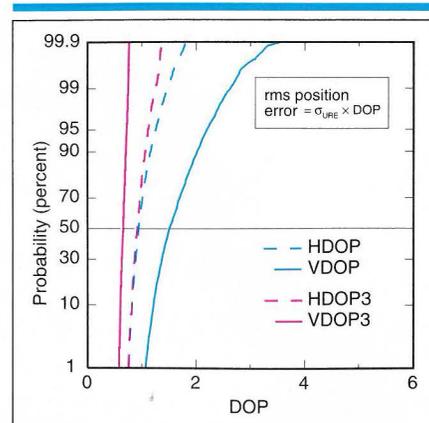


Figure 1. The cumulative probability distribution functions of DOP parameters with (HDOP3 and VDOP3) and without (HDOP and VDOP) a perfect receiver clock show that, for a given σ_{URE} , vertical position can be estimated much more accurately from a snapshot of GPS measurements if the clock bias is known *a priori*.

errors in the corresponding clock bias estimates (in units of length). The scatter shows no apparent correlation between the two errors. By contrast, Figure 2b, a plot of the errors in estimates of vertical position and clock bias, shows a strong linear correlation. Clearly, error in the clock bias estimate from a snapshot of pseudorange measurements is a good predictor of the error in the corresponding vertical position estimate. This observation serves as the basis of new algorithms for navigation and receiver autonomous integrity monitoring (RAIM) for precision approaches of aircraft to airport runways.

The observed correlation between the errors in the estimates of clock bias and vertical position is not a surprise. A change in receiver clock bias changes all pseudoranges equally: *all* increase, or decrease, by a common amount. The net effect of a change in user altitude is similar: *All* pseudoranges increase, or decrease, but, in general, not equally. Hence the correlation. Note, however, that if all the satellites were at the same elevation angle, a change in user altitude would indeed change all pseudoranges equally, and we wouldn't be able to distinguish between the user altitude and the clock bias.

How can we benefit from this observed correlation? Obviously, if we were to obtain somehow a "good" estimate of the receiver clock bias, we would know the error in its current snapshot-based estimate, and we could correct the corresponding vertical position estimate. Equivalently, given an accurate

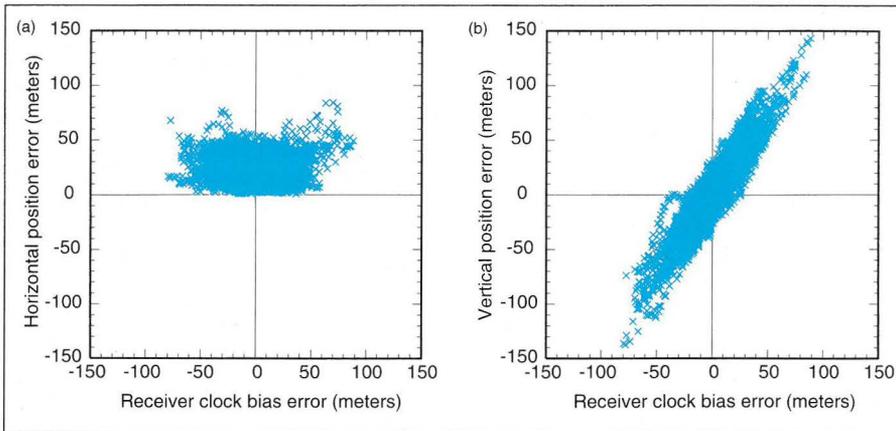


Figure 2. These plots show the scatter of errors in snapshot-by-snapshot estimates of position and clock bias. Errors in the estimates of (a) horizontal position and clock bias appear uncorrelated; error in the estimates of (b) vertical position and clock bias are highly correlated.

estimate of the receiver clock bias, we may reformulate the problem as one of 3-D estimation with an accompanying improvement in the vertical position estimate. These results basically confirm and explain the results of the previous section.

In general, we cannot expect to know the

true receiver clock bias. But we can estimate it from GPS measurements over a time period consistent with the stability characteristics of the clock. We can conclude from Figure 2 that the degree of improvement in the vertical position estimate would depend upon the quality of the clock bias prediction.

The horizontal position estimate would be substantially unaffected. So, the obvious questions are: How do we estimate receiver clock bias from the measurements so as to transform the 4-D problem into a 3-D problem? and What is required of the clock in order for this approach to work?

CLOCK MODELING

It is well known that the rms error in clock bias estimate (in units of length) based upon a single snapshot of pseudorange measurements is:

$$\sigma_b = \sigma_{URE} \times TDOP$$

where σ_{URE} , introduced previously, is the rms error in the pseudorange measurements, and $TDOP$ is the time dilution-of-precision parameter reflecting the satellite geometry. In the presence of selective availability, σ_{URE} has been estimated conservatively as 33 meters (the observed value is about 25 meters). For the fully operating GPS constellation, $TDOP$ usually ranges between 0.75 and 1.25. Taking a typical value of 1 for $TDOP$, we can easily calculate that the rms error in clock bias estimated from a



Sponsored by
The Institute for International Research

TRACKING CARGO BY LINKING GPS WITH COMMUNICATION NETWORKS

May 30-31, 1996
Watergate Hotel • Washington, D.C.

FEATURING PRESENTATIONS FROM:

- ◆ Trimble Corporation
- ◆ Rockwell Int'l
- ◆ Mark VII Logistics Services, Inc.
- ◆ Orbcomm
- ◆ The Pentagon
- ◆ Highway Master
- ◆ UPS Worldwide Logistics
- ◆ American Mobile Satellite Corporation
- ◆ Canadian Marconi
- ◆ Werner Enterprises
- ◆ ADT Auto Auction Transport

Post-Conference interactive workshop on
Friday, May 31 from 1:30 - 5:00.

Process Development & Re-engineering for Integration of GPS

For more information, call Beth Debra Kallman, Vice President
Logistics Division at 212-661-3500, Ext. 3820, or fax 212-599-2192.



VEHICLE NAVIGATION LOCATION AND CONTROL

18-19 JUNE 1996
LONDON

Building on earlier successful events, the British and German Institutes of Navigation are holding their third joint European conference at Church House, Westminster. The conference will focus on the significant developments taking place in the Advanced Transport Telematics (ATT) sector worldwide; there will be an associated exhibition.

Sessions will include:

- intelligent transport systems*
- public transport and fleet management*
- the intelligent highway*
- positioning technologies*
- multi-modal influences*

The conference will represent the European focus of the state of the art in this burgeoning topic.

For further information please contact:

The Royal Institute of Navigation
1 Kensington Gore
London SW7 2AT
Telephone: +44 (0)171 589 5021
Facsimile: +44 (0)171 823 8671
Email: rindir@atlas.co.uk



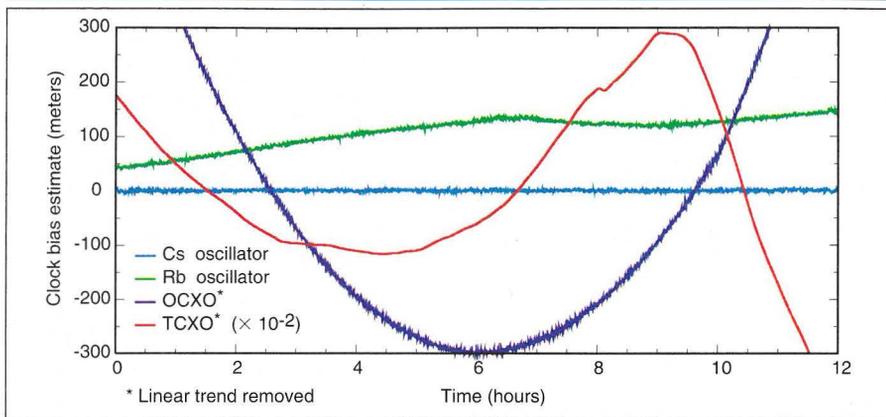


Figure 3. The stability of different kinds of clocks varies widely. Shown here are the changes in clock bias of several receiver clocks relative to GPS Time, estimated from laboratory measurements.

single snapshot of pseudorange measurements is $\sigma_b \approx \sigma_{URE} \approx 33$ meters. In differential mode, TDOP is unchanged but the measurement error is significantly lower, say, $\sigma_{URE} \approx 2$ meters. The corresponding rms error in clock bias estimate is then about 2 meters.

If we had available GPS measurements over a time interval (t_0, t_1) during which the

frequency drift rate of the receiver clock is stable, we could model the clock bias at time t simply as a quadratic function

$$b(t) = b_0 + b_1(t-t_0) + b_2(t-t_0)^2$$

and estimate parameters b_0 , b_1 , and b_2 from the available measurements. Given k statistically independent measurement snap-

shots, the rms error in the clock bias estimate will be

$$\sigma_b \approx \frac{\sigma_{URE}}{\sqrt{(k-3)}}$$

Given snapshots of differentially corrected pseudorange measurements over a 30-minute interval and assuming a correlation time of 3 minutes (so that we can assume we have 10 independent measurements), then

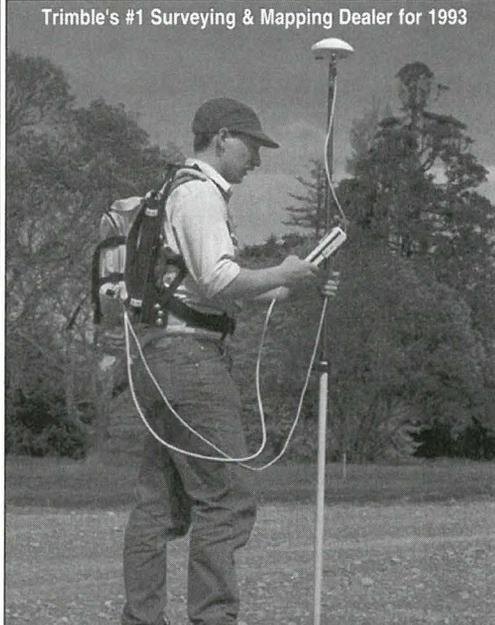
$$\begin{aligned} \sigma_b &\approx \frac{\sigma_{URE}}{\sqrt{(k-3)}} \\ &\approx \frac{2}{\sqrt{(10-3)}} \\ &\approx 0.8 \text{ meters} \end{aligned}$$

If this model were valid for a receiver clock, we would be able to predict clock bias with an rms error of about 1 meter, or equivalently, 3 nanoseconds. We will see next how well these simple calculations check out with actual measurements.

Figure 3 shows the behavior of several kinds of receiver clocks as observed in our laboratory over a 12-hour period. We included clocks using an atomic cesium beam oscillator, an atomic rubidium vapor-cell oscillator, an oven-controlled crystal oscillator (OCXO), and a temperature-compensated crystal oscillator (TCXO). The figure shows 3-second samples of clock bias estimates in meters for each clock based on GPS measurements with differential GPS (DGPS) corrections. The clock bias estimates were computed using the usual 4-D estimation from the DGPS data, snapshot by snapshot. The plots are intended to be representative of the typical performance of each kind of clock. Our purpose is to draw attention to certain qualitative features of the different clocks rather than to provide specific quantitative results.

In Figure 3, the excellent frequency stability of the cesium oscillator (with a price tag of about \$20,000) is obvious: no frequency offset or drift is apparent. The rubidium oscillator (costing from \$2,000 to \$4,000) also shows no significant frequency offset, but signs of a slight drift are clear. Such drift, however, appears constant over 30–60 minutes. The OCXO (costing about \$1,000) had accumulated a small frequency offset (about 1 part in 10^8) so that during our test the clock bias actually changed by about 2 meters per second on average. We have taken out the effect of this offset via linear regression and have plotted the residuals, which show a clear quadratic trend. The behavior of both atomic clocks and the OCXO appears consis-

Trimble's #1 Surveying & Mapping Dealer for 1993



GeoPlane Services
GPS Sales • Training • Equipment Rental



A Division of Universal Enesco, Inc.

(800) 229-3114
(713) 789-7296 Outside USA
Fax: (713) 977-9361

Certified Dealer for:

 **Trimble**
OMNISTAR
GLB Electronics Inc.

DCI ↔ SECO

Our staff of professionals have the proven experience and expertise to help you cost effectively meet your needs with GPS technology.

GPS Sales
GPS Rentals
GPS Consulting

GIS Consulting
Certified Repair
Certified Training

1811 Bering Drive, Suite 200, Houston, Texas 77057

tent with the proposed model: Clock bias is indeed predictable, given recent data. The TCXO (costing about \$100 or less and common in many of the low-cost GPS navigation receivers) had accumulated a frequency offset of about 1 part in 10^6 , and its bias changed by about 200 meters per second on average. Again, the residuals from a linear regression are plotted. These residuals, however, are two orders of magnitude larger than those for the OCXO and swing widely and wildly. It's clear that the TCXO does not exhibit the requisite stability of frequency drift rate, and the proposed model doesn't fit. Experience with several other TCXOs was similar.

How well does the proposed clock model fit the atomic clocks and the OCXO? For the OCXO, the answer is given in Figure 4; as might be expected, the atomic clocks did better. Figure 4 gives the postfit residuals for the OCXO bias estimates shown in Figure 3. At each epoch, a quadratic model was fitted to the clock bias estimates obtained over the previous 30 minutes, and clock parameters b_0 , b_1 , and b_2 were estimated. The residual at each epoch is defined as the discrepancy between the predicted value of clock bias based on the fitted model and the actual mea-

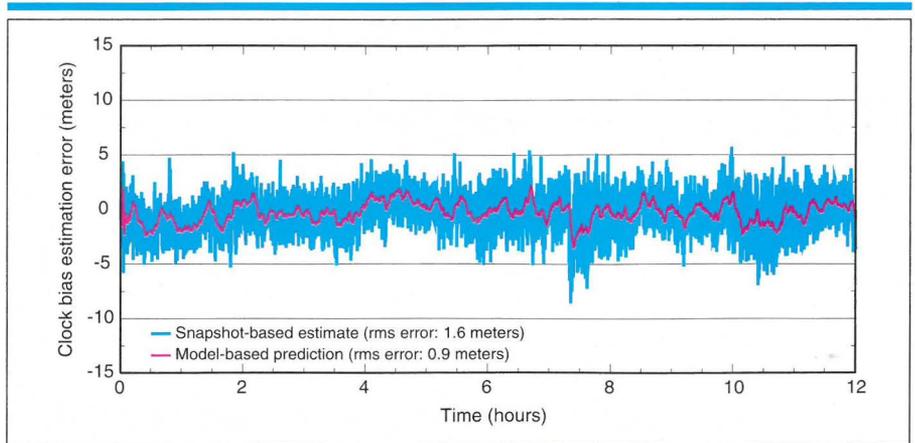


Figure 4. This plot shows the error in modeling a receiver OCXO clock based on GPS measurements over a moving 30-minute window. The rms clock bias prediction error is about 3 nanoseconds, or in terms of the distance traveled by a radio wave in such a time interval, about 1 meter.

surement. The rms error in the snapshot-based clock bias estimates is 1.6 meters; the rms error in model-based clock bias predictions is less than 1 meter. Both rms values are consistent with the results of our simple calculations given earlier in this section. Given that a receiver clock is predictable with an

rms error of 1 meter, according to Figure 2, a significant improvement in the quality of vertical position estimates can be effected through 3-D estimation.

CLOCK-AIDED NAVIGATION

The data presented so far were taken in the

BLILEY QUARTZ CRYSTALS



Give Enhanced Precision to Your GPS Receivers/ Navigational Guidance Systems

Here's how our BK3 Series performs

Frequency:	16.368327 MHz
Crystal Cut/Mode:	AT/Fundamental
Initial Accuracy:	+5/-10 ppm @ +25°C
Tolerance (vs Temp.):	±15 ppm (-30°C to +80°C)
Aging:	1 ppm/year @ +25°C
Resistance:	15 ohms maximum
Internal Mount:	Three Point (shock stable)
Holder Style:	Low Profile HC-35/U

And, our Low Profile BK3 Series crystals can be factory designed to your unique specifications. Call us - we can help you!



BLILEY ELECTRIC COMPANY

2545 West Grandview Blvd.
P.O. Box 3428, Erie, PA 16508
(814) 838-3571 • Fax: (814) 833-2712

INNOVATIVE ANTENNA SOLUTIONS
From Allis Communications Company/AAT



High Performance Broadband, Narrowband, and Multiuse Antennas



For Today's Positioning & Wireless Communications Applications: GPS/GLONASS, Cellular, PCS and SATCOM



Is your application in need of an antenna fix? Are your antennas giving you a headache? Let RF-Doc™ Antennas perform system surgery! Call for an appointment today...

U.S.A. Advanced Antenna Technology, Inc. San Diego: Tel:(619)737-8128 Fax(619)741-8898
TAIWAN Allis Communications Company LTD. Taipei: Tel:(02)695-2378 Fax(02)695-7078

benign environment of a laboratory. How would the rubidium oscillator and the OEXO, which were found to have the requisite stability in the laboratory, hold up in an aircraft environment that is characterized by vibrations, accelerations, and rapid changes in temperature and pressure? We have conducted several flight tests to answer this question.

In the flight tests, a mobile GPS data-collection system was placed in a rack in a Beechcraft King Air 200 — a small twin-engine airplane. The rubidium oscillator sat next to the receiver, with no provision to isolate it from any environmental stresses. The GPS measurements were recorded as the aircraft executed a set of approaches. Data were also recorded simultaneously at a local, ground-based GPS reference station. Both data sets were processed postflight to generate position estimates in both 4-D (snapshot-based) and 3-D (clock-aided) mode. The “ground truth” for each flight was obtained by postprocessing carrier-phase measurements recorded independently by a pair of GPS receivers, one aboard the airplane, the other at a surveyed location on the ground.

Figure 5 shows the benefits of clock aiding in two of the flight tests. Shown are errors in vertical position estimates obtained from 4-D estimation snapshot by snapshot, and those from 3-D estimation where we have used an estimate of the clock bias based on measurements over the previous half-hour. Clearly, clock modeling offers a significant improvement. The snapshot-based DGPS vertical position estimates in both flight tests show a bias of 0.6 meters, attributed to discrepancies in tropospheric propagation delay models. As expected, this bias carries over to the clock-aided estimates. The standard deviation of the error for the snapshot-based vertical position estimates is 1.6 meters in each test; the corresponding value for clock-aided navigation is 0.9 meters. Another noteworthy feature of our results is that clock aiding consistently avoids the peak errors obtained with 4-D estimation. This is important in view of integrity monitoring requirements in civil aviation where a large, undetected excursion from the defined flight path caused by navigation error may be hazardous.

ADDITIONAL BENEFITS

A predictable receiver clock can further benefit integrity monitoring and also aid in the processing of carrier-phase data.

RAIM. The objective of RAIM is to ensure that error in a position estimate is not so large as to compromise the user’s safety. For user requirements no more stringent than those for

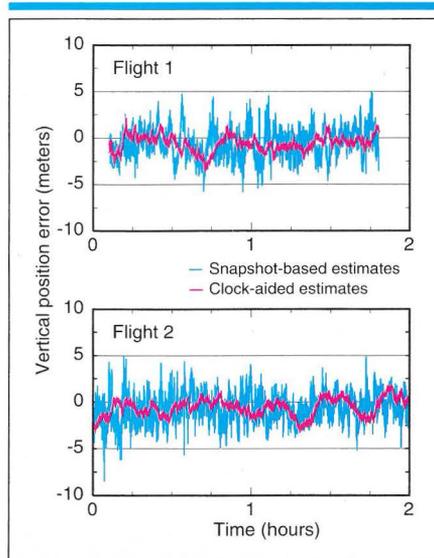


Figure 5. These results of two clock-aided navigation in-flight tests show that vertical position error can be reduced significantly with clock modeling and prediction.

nonprecision aircraft approaches (cross-track accuracy at 95 percent probability of 0.3 nautical miles or 0.56 kilometers), RAIM is essentially equivalent to ensuring that there are no significant measurement anomalies. Faulty measurements must be detected and excluded from position estimation. A predictable receiver clock provides the user with a built-in test for fault detection and exclusion. The premise is that any satellite system anomalies affecting the user position estimate would affect the clock bias estimate too. A test would consist of 4-D estimation snapshot by snapshot. If the clock bias so estimated deviates from the correct answer by more than can be explained by user range error and TDOP, there’s a strong likelihood that one or more of the measurements is faulty.

Carrier-Phase Processing. Two important problems related to carrier-phase processing are detection of cycle slips and resolution of integer ambiguities. With a perfect clock, the detection and repair of cycle slips would be straightforward. A perfect clock would also allow us to resolve integer ambiguities within the finer structure of single differences of phase measurements at the reference and mobile receivers, rather than the usual double differences. As in our earlier discussion, a predictable receiver clock can offer a significant leverage in solving both problems. Obviously, the accuracy of prediction will have to be commensurate with the wavelengths of signals being processed.

CONCLUSIONS

A receiver clock with predictable behavior can offer several important benefits to the user. In practice, the minimum required of such a clock appears to be that its frequency drift rate be stable for a long enough period for the model parameters to be estimated from satellite measurements. This approach, referred to as *clock aiding*, has been found to be practical. Widespread use, however, must await development of robust clocks at lower prices.

ACKNOWLEDGMENTS

The work in this article was sponsored by the Federal Aviation Administration (FAA) under Interagency Agreement DTFA01-95-Z-02046. The author is grateful to Joseph F. Dorfler and Michael E. Shaw of the FAA Satellite Program Office for their support. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by FAA. ■

MANUFACTURERS

The GPS data used for the analyses were obtained from a NovAtel GPSCard, manufactured by **NovAtel Communications, Ltd.** (Calgary, Alberta, Canada).

This article is based on the paper “Adaptive Modeling of Receiver Clock for Meter-Level DGPS Vertical Positioning,” presented at ION GPS-95, held in Palm Springs, California, in September 1995.

For reprints (250 minimum), contact Mary Clark, Marketing Services, (541) 343-1200.

Further Reading

For a basic introduction to the use of clocks in GPS, see

- “Time, Clocks, and GPS” by Richard B. Langley in *GPS World*, November/December 1991, pp. 38–42.

For a discussion of clock modeling in a Kalman filter implementation of a navigation algorithm, see

- “GPS Navigation Algorithms” by P. Axelrad and R.G. Brown, Chapter 9 in *Global Positioning System: Theory and Applications*, Vol. I, edited by B.W. Parkinson and J.J. Spilker, Jr., published as Vol. 163 of Progress in Astronautics and Aeronautics, American Institute of Aeronautics and Astronautics, Inc., Washington, D.C., 1996, pp. 409–433.

For a discussion of RAIM, see

- “Receiver Autonomous Integrity Monitoring” by R.G. Brown; Chapter 5 in *Global Positioning System: Theory and Applications*, Vol. II, published as Vol. 164 of Progress in Astronautics and Aeronautics, pp. 143–165.