Software-based Multipath Mitigation: Sampling for Multipath Invariance

R. Eric Phelts, Jonathan Stone, Per Enge, J. David Powell

Stanford University, Stanford, California

BIOGRAPHY

R. Eric Phelts is a Ph.D. candidate 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. in Mechanical Engineering from Stanford University in 1997. His research involves multipath mitigation and passive ranging applications.

Jonathan M. Stone is a Ph.D. candidate in the Department of Aeronautics and Astronautics at Stanford University. He received his B.S. in Aerospace Engineering from Texas A&M University in 1991, and his M.S. at Stanford University in 1992. His research focuses on providing precise position information using GPS satellites and pseudolites.

Per Enge is a Professor of Aeronautics and Astronautics at Stanford University, where he has been on the faculty since 1992. His research deals with differential operation of GPS for landing aircraft. Previously, he was an Associate Professor of Electrical Engineering at Worchester Polytechnic Institute.

J. David Powell is a Professor of Aeronautics and Astronautics at Stanford University, where he has been on the faculty for 26 years. His research deals with land, air and space vehicle navigation and control. He is the co-author of two control system textbooks.

ABSTRACT

Multipath potentially introduces a significant bias error into GPS code phase positioning solutions. That error is dependent on the amplitude, delay, phase, and phase rate of the multipath with respect to the line-of-sight. Many current mitigation techniques attempt to eliminate or reduce the impact of this error source. These techniques vary as widely in hardware and software complexity as they do in overall mitigation performance. All, however, perform very differently for different multipath conditions. They exhibit a characteristic dependence on the multipath parameters. For example, most will easily mitigate long-delay multipath, however, few or none perform as well for multipath at short relative delays.

In this paper, a new software-based mitigation approach is introduced. It is asserted that since the C/A codes are not orthogonal, for each satellite autocorrelation function, a different location which is off the main lobe will remain virtually unaffected by multipath. A method is given for sampling at and around the autocorrelation function peak and/or sidelobes with fidelity, such that the tracking errors may become essentially multipath invariant-- independent of the multipath parameters. Theoretical performance plots are shown for

validation. It is asserted that since the remaining error bias is essentially constant, it may subsequently be modeled, calibrated, and/or estimated away.

An algorithm for finding this point is provided. It was implemented in software on a Mitel Semiconductor receiver. Preliminary analyses and results are given from data taken in real-time with both a GPS signal simulator and actual satellite data. The experimental data supports the claim that there are multipath invariant points corresponding to each PRN code and that with some iteration and smoothing, they can be found. It is believed that with more refinement and proper implementation, this approach may provide superior mitigation performance to many techniques currently in use.

INTRODUCTION

Multipath Effects

For GPS users, multipath (MP) is caused by reflections of the satellite signal from the ground or from nearby buildings or other obstacles. Multipath errors result when the receiver receives the direct or line-of-sight (LOS) satellite signal via multiple paths and processes the combined signal as if it were a single path. These errors are particularly difficult to remove since, in general, the following apply:

1) Multipath is not spatially correlated. MP signals affecting a receiver at one location will not affect a receiver at another location in the same way.
2) Multipath errors are not zero mean. This is particularly true for MP signals with relatively large amplitudes. Consequently, even infinite smoothing of the pseudorange cannot guarantee unbiased position errors [Van Nee, 1992].
3) The pseudorange measurement is derived from a code-tracking delay-lock loop (DLL). DLL's essentially attempt to derive time-of-arrival measurements from measurements of incoming signal amplitudes. This is accomplished by maximizing the signal code autocorrelation function. In the receiver, this translates to using a minimum of 2 correlators to straddle the peak (e.g., such that Early-Late=0). Since the combined LOS and MP autocorrelation function will have a distorted shape, the distortion introduces a tracking error into the DLL. An example of ideal and MP-corrupted autocorrelation functions is given in Figure 1.
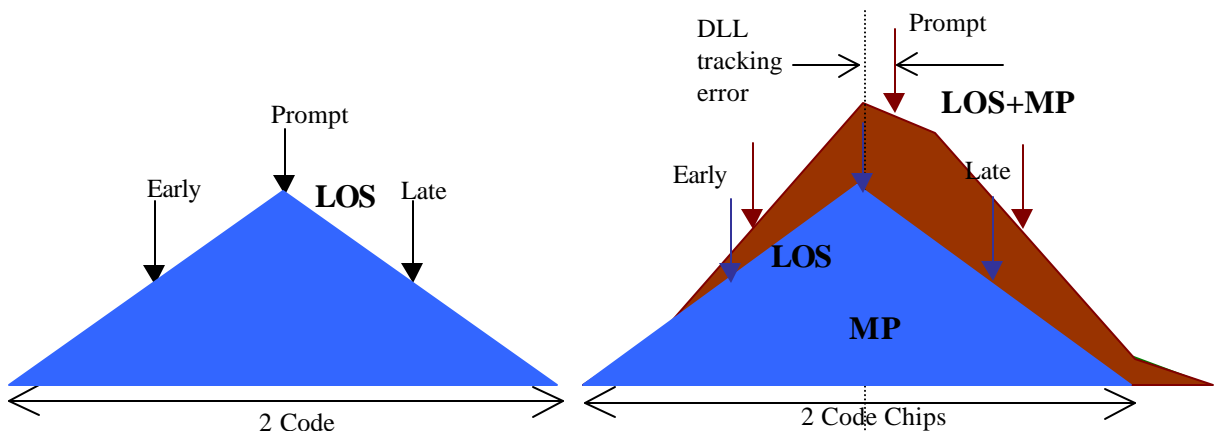
Figure 1. Ideal autocorrelation peaks with and without multipath (pictured in-phase)
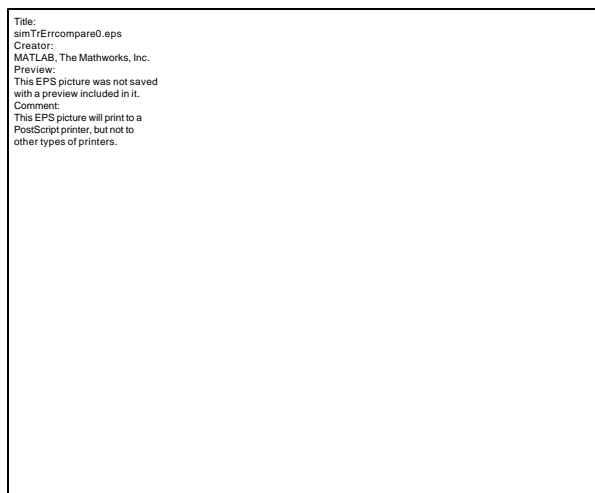
4) Pseudorange errors due to multipath, in general are nonlinear functions of MP amplitude delay, phase and phase rate [Parkinson, 1996]. Accordingly, changes in any of these parameters may significantly change the tracking response of the DLL.

A plot of position error vs. relative MP delay is shown above in Figure 2 for the (noiseless) case of a receiver having a precorrelation bandwidth of 2MHz. For comparison, error envelopes are given for a correlator chip spacing of both 0.5 and 0.1-chips of the L1 C/A code. As expected, because of the narrow bandwidth being simulated here (causing a flattening of the autocorrelation peak), the tracking loop with the narrow correlator chip spacing does not significantly outperform the wider one.

Current Mitigation Techniques

A great many multipath mitigation techniques have been developed over the years. Most of these share at least two things in common. First, they concentrate almost exclusively on extracting information from the peak of the autocorrelation function. This is intuitive since most of the signal power is concentrated here, and for orthogonal code sequences, the main lobe theoretically contains all the information of the signal. Although the C/A codes are not orthogonal, the sidelobes of the autocorrelation function are or the most part ignored. Second, they all tend to perform well against long-delay multipath, with relatively strong signal-to-noise ratios (SNR). However, when faced with low-SNR multipath and/or multipath with relatively short delays (e.g., less than 30 meters) their performance degrades to that of a cnventional, 1-chip spacing DLL.

Most current MP mitigation can be grouped into two major categories: *Separation* and *Estimation*. Separation techniques (e.g., the Narrow Correlator [Dierendonck et al., 1997], Strobe Correlator [Garin et al., 1997], etc.) essentially attempt to separate the LOS and MP signals. These approaches attempt to track only the LOS signal and thereby reduce or eliminate the effects of the multipath. Estimation techniques (e.g., MEDLL [Townsend et al., 1995], MET [Townsend et al., 1994], MRDLL [Gadallah et al., 1998], etc.) attempt to estimate the parameters of the LOS and/or MP signals, and approximate their combined effect on the tracking errors. They may also or use some combination separation and estimation schemes to form a correction factor for (or rather an estimate of) the code tracking error [Cahn et al., 1997].

Title:
simTrErrcompare0.eps
Creator:
MATLAB, The Mathworks, Inc.
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
PostScript printer, but not to
other types of printers.

Figure 2. Tracking error vs. MP relative delay.

Both classes of techniques rely on an ability to somehow distinguish the multipath from the line-of-sight. This is most readily done by special signal processing of the autocorrelation functions and/or discriminator curves. Extra hardware (i.e. more than the usual Early, Late and sometimes Prompt correlators) and often a wider bandwidth, is frequently employed for this purpose. A fundamental limitation these methods must overcome is their sensitivity to the changing characteristics of the multipath. For example, the closer the MP parameters match those of the LOS, the more difficult it becomes to either separate or estimate one from the other. This explains the characteristic degradation in performance these techniques suffer when the MP relative delays are very short.
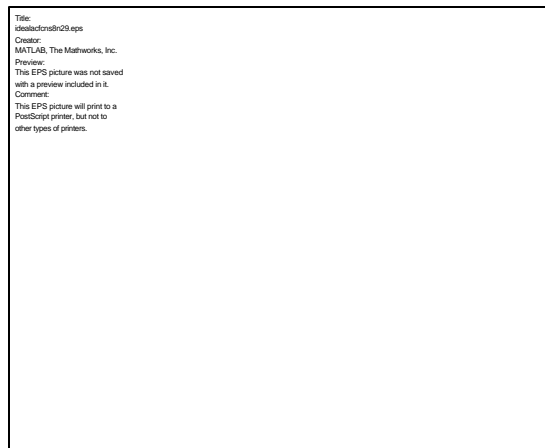
The E1-E2 Tracker, however is an example of a mitigation technique that differs from most. It controls the receiver correlators to obtain samples of the earliest, relatively "undelayed" portions of the main lobe of the autocorrelation function [Mattos, 1992]. The difference between the sampled slope and ideal, expected slope of this low-SNR region of the main lobe is used to dynamically adjust the DLL. This is a separation and estimation based technique. However, like many others, it ignores the effect of multipath from sidelobes at long relative MP delays. For best performance, it also requires a maximum receiver precorrelation bandwidth (e.g., 20MHz) for accurate estimates of the slope of the leading edge of the primary peak. It performs best against medium and long-delay multipath only. Also, because it uses samples of the main lobe at low correlation values it has received a somewhat dishonorable mention because of its poor noise performance [Dierendonck et al, 1997].

MITIGATION BY MULTIPATH INVARIANCE

Concept

One possible way to avoid the characteristic dependence on multipath parameters is to reference a location on the autocorrelation function that does not change as a function of the multipath amplitude delay, phase or frequency. It is clear from Figures 1 and 2 that if such a point exists, it does not in general lie at or even close to the autocorrelation peak. It is proposed that for each PRN code autocorrelation function there exists a number of points that, sampled, averaged, and referenced properly in a position solution may produce tracking errors that are virtually *multipath invariant*.

Figure 3 below shows the autocorrelation function for PRN 8 with a few neighboring sidelobes. Because the sidelobes nearest the peak for PRN 8 are adjacent to the main lobe,

Title:
idealacfcns8n29.eps
Creator:
MATLAB, The Mathworks, Inc.
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
PostScript printer, but not to
other types of printers.

Figure 3. Normalized ideal autocorrelation functions.

DLL's tracking this PRN will have one of the worst MP performance curves—since the tracking error will be non-zero for relative MP delays of up to 3.5 chips. Conventionally, it is assumed that the tracking error is zero for MP relative delays greater than 1.5 chips. When multipath is present on this PRN code (and on many others), however, the presence of sidelobes or regions with non-zero slope will interfere with the DLL performance just as would a small replica of the peak.

Figure 3 also shows the autocorrelation function for PRN 29, note that it's sidelobes are close to zero (have approximately zero amplitude and have zero slope) "earlier than" (to the left of) the main lobe. This means that the DLL will be virtually unaffected by multipath at relative delays greater than 1.5 chips. Note that this is irrespective of the number of multipath signals and the characteristics (amplitude, phase, and phase rate) of the multipath at a relative delay greater than 1.5 chips.

Indeed if it were possible to reliably track the sidelobe (assuming positive magnitude) of PRN 8 located at –5 chips earlier than the peak with a conventional DLL, it too would have the MP tracking error characteristics of SV 29. Multipath having relative delays greater than 1.5 chips, and less than the distance to the next nearest sidelobe, would be ignored. In other words, for some finite span of delays, *plateaus* in the autocorrelation function remain multipath invariant. That is, their amplitudes increase or decrease by almost negligible amounts. More significantly, however, they have zero slope. Superimposing these small, constant amplitude offsets anywhere on the autocorrelation function does not alter the shape (hence tracking error) of the function at that location.

Assumptions

It is reasonable to assume that multipath having relative delays of say three chips (approximately 880 meters) or more are unlikely and/or have extremely low mean SNR's compared to the LOS autocorrelation peak and sidelobes. The attenuation assumption is most adequate when assuming changing receiver-reflector geometry. Still, the fading frequency tends to be higher at greater distances between the reflector and the receiver [Van Nee, 1992]. Both of these factors tend to "whiten" any spurious MP interference and bring it closer to zero mean. Consequently, it filters out more easily.

Additionally, in order to locate and make use of this multipath invariant region we further assume:
1) The relative distance between the main lobe and each respective sidelobe remains constant. In other words, the C/A code may be either delayed or advanced as a whole (e.g., by the ionosphere), but the chip lengths remain constant.
2) The receiver is tracking a given satellite such that it's PRN (hence its autocorrelation function) is known.
3) Thermal noises, cross-correlation noises due to the presence of other visible GPS satellites, and multipath are the only interference sources present.
4) The relative positions of the receiver correlators are controllable in software.

Implementation

Infinite Bandwidth

Figure 5 illustrates one possible way to make use of this multipath invariant (MPI) region. First, we locate the longest (e.g., greater than 3 chips) plateaus to the left of the main lobe and or sidelobes. Then we observe (in the infinite bandwidth case) that the most
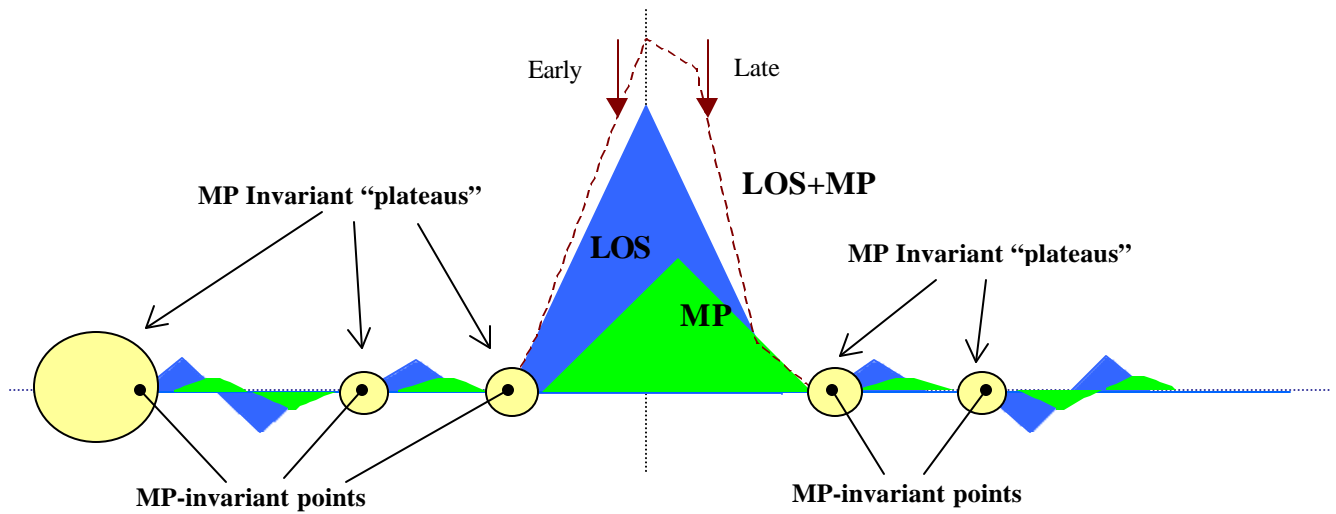
Figure 5. Illustration of an autocorrelation function and the nearest MP invariant regions.

definable characteristics of this region occur at the corners or slope transition points alongside the early slopes of the side and main lobes. If we can control the relative positions of the correlators, we can then iteratively search for this location. Since the ideal distance from the main lobe for a given PRN is known, it may be possible to correct our MP-biased tracking loop (DLL) solution based on the *relative* distance measured from this MP-invariant point. Figure 6 shows how this correction might be applied in the ideal (noiseless, infinite bandwidth) case.

A simple algorithm for seeking out and locating the MPI point follows:

1) Generate (offline) the ideal autocorrelation functions corresponding to each PRN and compose a look-up table corresponding to the ideal distance between the peak and the MPI point.
2) Once the receiver is tracking a satellite in a given channel, use knowledge of the ideal distance to the MPI location for the corresponding PRN and reasonable bounds on the current tracking error to bound the desired point. (See Figure 7.)
3) Perform a rotation (in software) such that the MPI location becomes a minimum. (See Figure 7.)
4) Optimize (minimize) the localized region of the autocorrelation function to find the location of the MPI point with respect to the primary (tracking) correlators.
5) Once the point is identified, correct the tracking loop solution by the difference between the Measured and Corrected DLL positions. (See Figure 6.)
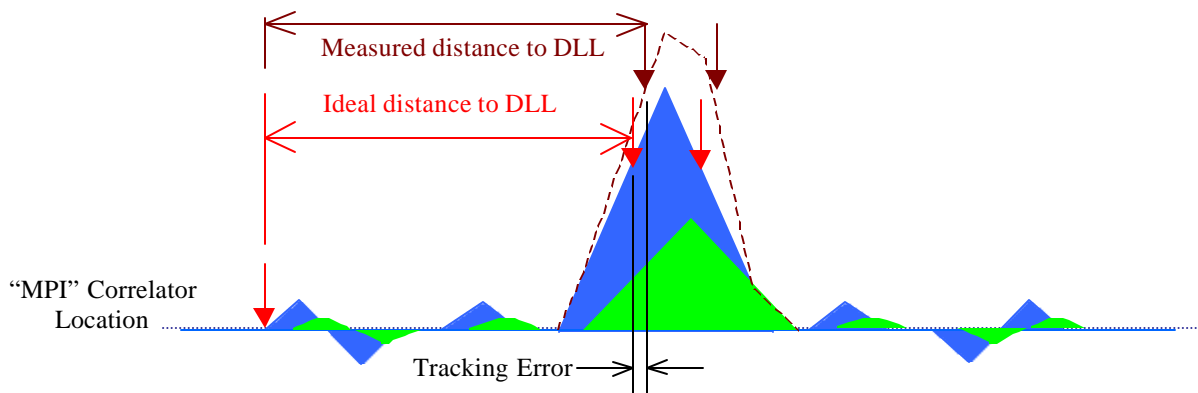
Figure 6. Updating the primary DLL tracking solution with the corrected, relative position to the MPI point.
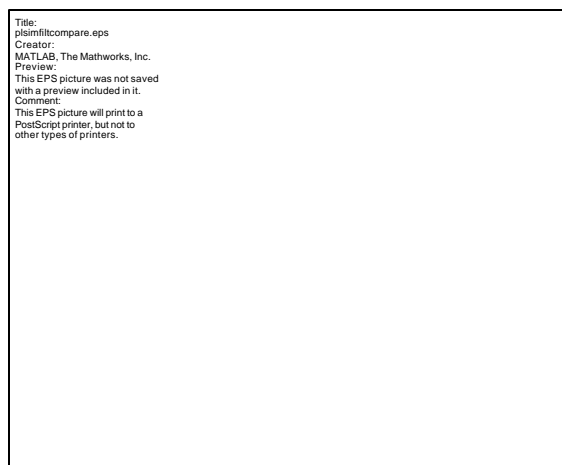
$$C'_{MPI} = R \cdot \begin{bmatrix} d \\ C_{MPI} \end{bmatrix}$$

* C is a function of SNR
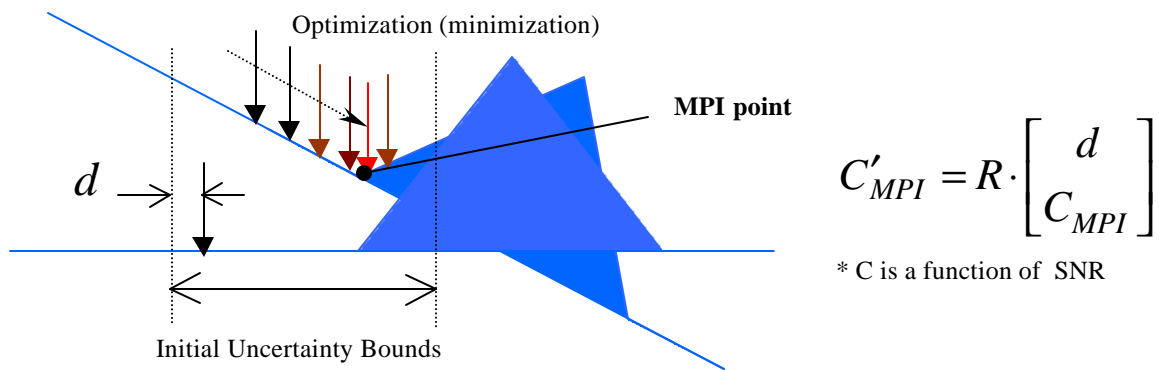
Figure 7. Illustration of optimization of the MPI region to locate the MPI point.
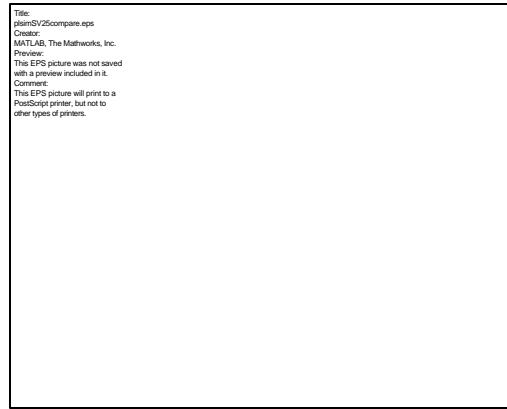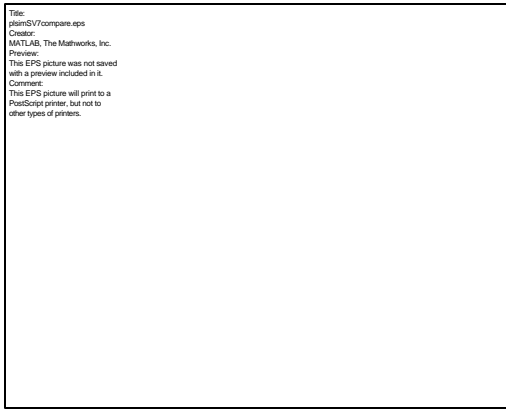
Finite Bandwidth

Because GPS receivers have finite bandwidth, the sharp peaks and corners of the ideal autocorrelation functions become more smooth and rounded. Recall that the error envelopes of Figure 2 corresponded to a receiver precorrelation bandwidth of 2MHz. This bandwidth limitation causes a significant rounding of the features of the (simulated) autocorrelation function. In attempt to verify how well the simulation modeled the bandlimiting effects of the receiver, the peak of a main lobe generated in simulation was compared to actual samples of an autocorrelation function taken from the (2MHz-bandlwidth) Mitel receiver.

Figure 8 shows a comparison of the resulting two curves. The maximum amplitude of the simulated peak was scaled to match that of the maximum sample from the receiver, however, no other adjustments were made. For the Mitel data a high-elevation satellite was tracked (to minimize multipath interference) and all twelve channels were used to track the satellite. The receiver has two correlators (Early and Late) per channel. Channel 1 was assigned as the "master" or primary channel and the others were "slaved" off of it. The DCO's (Digitally Controlled Oscillators) of the 11 slave channels were used to command the respective correlators to approximately equally-spaced offsets later than (to the right of) the Early correlator of the primary channel.

Similarly, the MPI point will be affected by finite bandwidth. The early sides of the

Figure 8. Comparison of simulated and experimental 2MHz autocorrelation peaks.

Title:
plsimSV7compare.eps
Creator:
MATLAB, The Mathworks, Inc.
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
PostScript printer, but not to
other types of printers.

Title:
plsimSV25compare.eps
Creator:
MATLAB, The Mathworks, Inc.
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
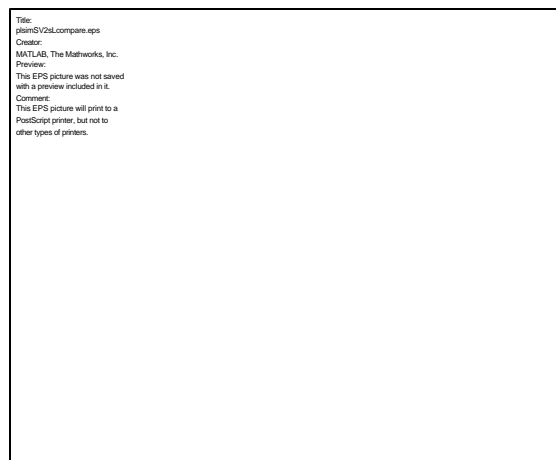PostScript printer, but not to
other types of printers.

Figure 9. Live satellite data showing the effects of finite bandwidth and (filtered) thermal noise on samples of the autocorrelation peaks and sidelobes.

autocorrelation functions of PRN 7 and PRN 25 are shown in Figure 9. Figure 10 shows an enlargement of the sidelobe at a proposed MPI region for PRN 2. For this PRN, the ideal MPI point would be located at –9 chips from the peak of the main lobe. However, due to finite bandwidth the actual MPI location may be actually be *slightly earlier than*—to the left of—the ideal. Still, because the optimization or search requires a rotation to be performed. As a result of this rotation (and due to noise concerns), the resulting "effective" MPI point may sometimes even be *later than*—to the right of—the ideal. Despite this fact, as long as its distance from the ideal MPI point is sufficiently small, this should not significantly affect the claim of MP invariance.

Finite Bandwidth and Thermal Noise

At present a rigorous analysis of the ability to locate the MPI point in the presence of thermal noise has not been done. Still the experimental evidence suggests that it is possible to iteratively search for and converge to this point. The most encouraging of such qualitative evidence is how easily the sidelobes can be distinguished from the noise floor. Figures 8-10

Title:
plsimSV2sLcompare.eps
Creator:
MATLAB, The Mathworks, Inc.
Preview:
This EPS picture was not saved
with a preview included in it.
Comment:
This EPS picture will print to a
PostScript printer, but not to
other types of printers.

Figure 10. Enlargement of the MPI region sidelobe for PRN 2.

were generated using live satellite data from a 12-channel Mitel receiver using the master-slave techniques described previously. The data was filtered to approximately 0.01 Hz. Note that although autocorrelation sidelobes are have signal powers attenuated on the order of 26 dB with respect to the peak, with a reasonable smoothing bandwidth, we are able to identify well-defined sidelobes down to the noise floor.

The MPI approach may, however, require sampling at extremely low signal power for a time. Despite this requirement, there are two primary reasons for a possible improvement in overall multipath mitigation ability. First, the MPI point is based on a relative distance that is not changing with time. This means that the actual quantity of interest is a constant and accordingly may be filtered with a very long time constant. Because the primary channel (or primary correlator pair) is actively tracking the signal dynamics, the relative dynamics of the samples about the MPI point can be made negligibly small. The details of how this is done is beyond the scope of this paper.

Second, MPI convergence occurs *independent of the tracking loop*. In other words, tracking performance is not degraded as long as the MPI point has not been found. The tracking loop solution may be improved, however, once the MPI point has been found to within some prescribed variance. The trade-off is that this approach will require some fixed initialization time to achieve desired convergence.


RESULTS

Simulation & Theoretical Performance

Figure 11 shows (again) the bandlimited (2MHz) error envelopes for a standard DLL (0.5 and 0.1-chip spacing). Additionally it shows the "envelope" for the 0.5-chip standard DLL corrected using the MPI approach. As was the case for Figure 2, a SMR of 3dB was used, and the effects of thermal noise were not simulated.

We observe that in the noiseless, bandlimited case, the MPI routine produces an essentially constant position error which is invariant with the changes in multipath amplitude, delay, and phase. As previously observed, however, the conventional DLL's, however, produce the conventional MP dependent "envelopes" of position error. The constant offset in position error, is attributed to the finite bandwidth of the receiver. It is dependent, however on the amount of rotation used in the optimization routine. Still, since the offset does not depend on the multipath parameters, it may be obtained once and stored. It can then be subtracted from any subsequent tracking errors resulting from implementation of the MPI routine.

Provided the same rotation operation is performed for each PRN of interest, one or more of the following three methods can be used to find the offset:
1) *Simulation*: A good software simulation may effectively model bandlimiting effects of a given receiver, the biases can be approximated by the ones found by implementing the routine in software.
2) *Calibration*: Using a GPS signal generator or simulator or pseudolite, the bias may be found directly. Since the receiver's conventional DLL will not be corrupted by multipath, the bias can be computed as simply the mean difference between the ideal (infinite bandwidth) distance and the measured distance to the MPI point.
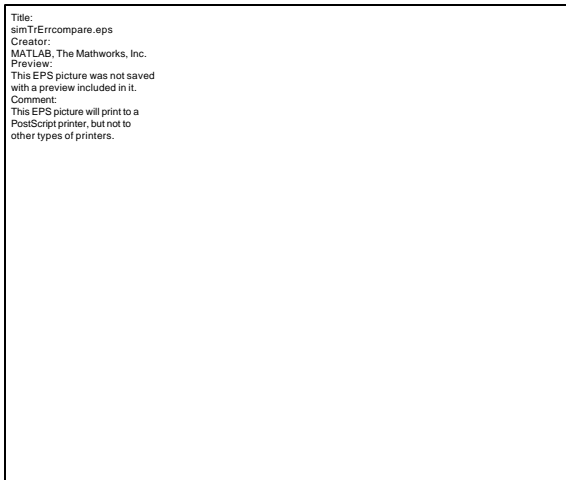
Figure 11. Comparison of simulated MPI and conventional DLL multipath error envelopes.

3) *Estimation*: If MPI locations for all the acquired PRN's are obtained at the early "corners" of similarly-oriented autocorrelation sidelobes, they will all have the finite-bandwidth bias. This is because the receiver front end filters all incoming signals the same. In this case, if the routine were implemented and corrections applied to all pseudoranges in the navigation solution, the bias (or at least some component on the biases) will appear as common-mode. Accordingly, that common-mode bias will become part of the clock bias term computed in the navigation solution.

Experimental Validation

In attempt to determine the bias between the actual MPI location (owing to finite receiver bandwidth) and the ideal, MPI approach was implemented using the Mitel receiver. A GPS signal generator was used to simulate the GPS satellite constellation for a static user. No multipath was introduced into the signals so the multipath affecting the code tracking loops was assumed to be negligibly small. Two master channels were used to track two satellites (SV25 and SV29). The ideal MPI points for PRN's 25 and 29 were located –5 and –26 chips respectively from the main lobe peak. In real time, two pair of slave channels (two secondary channels for each primary) were used to implement MPI algorithm. A Fibonacci optimization routine was used to minimize the MPI regions of the autocorrelation functions [Onwubiko, 1989]. As small a rotation angle as possible was experimentally determined and used to obtain an MPI location as close to the plateau (early) as possible. An interval of uncertainty of 1 chip (±0.5 chips) was used for each PRN.

Figure 12 shows a time trace of the iteration. After approximately 30-35 minutes, the trace for each satellite converges to a point earlier that the ideal corner by approximately 6 meters. The mean difference between the two traces is only 0.8 meters. This translates directly into the mean tracking error one could expect if this bias were used to modify the primary tracking loop solution.

The standard deviation was less than 2 meters. It should be noted, however that the variance of these traces reflect two random components. The first is the variance of the optimization (random convergence to the relative delay of the MPI location). The second is the variance of the primary code DLL. This is the reference from which the MPI location
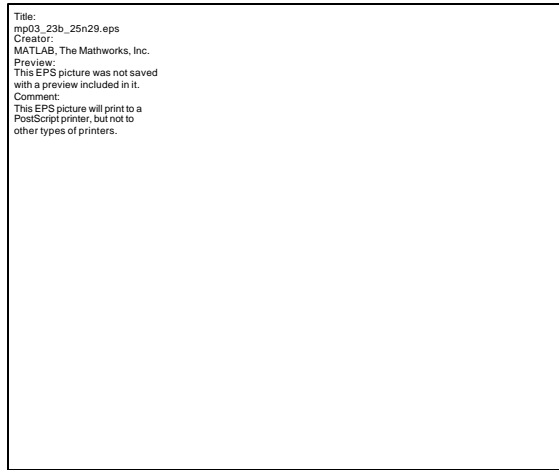
Figure 12. Convergence of MPI optimization.

must be measured. Hence, since code tracking loops normally have standard deviations on the order of 1-3 meters, we can conclude that the variance due to convergence alone is most likely significantly smaller than that of the difference reported here.

CONCLUSIONS

It is asserted that there is something useful to be gained from sampling the autocorrelation function at locations other than the peak. To this end, a new software-based approach to mitigating multipath was proposed. In particular, by sampling at the plateaus at the leading edges and bases of the autocorrelation function peaks or sidelobes, one may identify locations which may remain essentially invariant with changes in multipath amplitude, delay, phase, and frequency.

Results from a receiver code tracking loop simulation were generated and compared to actual receiver data for validation. In addition, a simple software algorithm was developed and implemented to locate MP invariant points. Once found they can be used to modify the code tracking loop solution. The (bandlimited, noiseless) simulation results clearly indicate that by successfully implementing this MPI approach, the theoretical tracking errors due to multipath are no longer envelopes which vary with the multipath parameters, but rather are constants. These constants may be found through calibration or modeled through simulation and subsequently removed

Using the Mitel Semiconductor receiver, data taken from actual satellites was used to show qualitatively that it is possible to extract useful information even at low relative signal power. The tradeoff is bandwidth. However, the relative distance to the MPI point is constant. Hence, relatively long filter time constants may be used to locate the point with fidelity. Once found it may be referenced to modify the tracking error of the primary DLL.

Using a GPS signal generator the MPI points for two different satellites were found to agree very closely. Since there was negligible multipath affecting these signals, this experiment was tantamount to calibration—to find the convergence bias corresponding to the chosen rotation angle. As expected, the time traces showed that the difference between the ideal distance and the measured distances to the MPI location for the two PRN's used, were approximately equal. In real time, both independent MPI channels converged to the offset. It is believed that this common-convergence condition should hold for each PRN provided that

the MPI regions of interest are bounded (on the right) by sidelobes and the same rotation operation is performed on the respective MPI regions.

FUTURE WORK

In order for this approach to become practical, the convergence times must be significantly decreased. In part this can be achieved by leveraging integrated carrier phase measurements to compute relative distances from a smoother reference than the primary code tracking DLL. Once accomplished, a true evaluation of the effectiveness of this approach could take place by first generating the experimental MPI tracking error envelopes. If the point is truly multipath invariant, the "envelope" will be virtually flat—as the simulation predicts. In addition, a rigorous quantitative noise analysis should be performed to determine the true performance limitations.

REFERENCES

1. Anon, 1997, *Autocorrelation Sidelobe Considerations in the Characterization of Multipath*, IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS, IEE Transactions on Aerospace and Electronic Systems, Vol. 33, No. 1, pp. 290-95.

2. Cahn, C. R., Chansarkar, M. M., 1997, *Multipath Corrections for a GPS Receiver*, Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION-GPS-97, Part 1(of 2), Proceedings of ION GPS-97, Vol. 1, pp. 551-57.

3. Dierendonck, A. J., Braasch, M. S., 1997, *Evaluation of GNSS Receiver Correlation Processing Techniques for Multipath and Noise Mitigation*, Navigation and Positioning in the New Age, Proceedings of the 1997 National Technical Meeting, Institute of Navigation, pp. 207-15.

4. Dierendonck, A. J., Fenton, P., Ford, T., 1992, *Theory and Performance of Narrow Correlator pacing in a GPS Receiver*, NAVIGATION, Journal of the Institute of Navigation, Vol. 39, No. 3, pp. 265-83.

5. Garin, L., Rousseau, J.-M., 1997, *Enhanced Strobe Correlator Multipath Mitigation for Code and Carrier*, Proceedings of the 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION-GPS-97, Part 1(of 2), Proceedings of ION-GPS, Vol. 1, pp. 559-68.

6. Gadallah, E.-S. A., Pachter, M., DeVilbiss, S. L., 1998, *Design of GPS Receiver Code and Carrier Tracking Loops for Multipath Mitigation*, Proceedings of the 11th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION-GPS-98, Part 1(of 2), Proceedings of ION GPS-98, Vol. 1, pp. 1041-54..

7. Mattos, P. G., 1996, *Multipath Elimination for the Low-Cost Consumer GPS*, Proceedings of the 1997 10th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GPS-96, Part 1(of 2), Proceedings of ION-GPS, Institute of Navigation, Vol. 1, pp. 665-71.

8. Onwubiko, C. O., 1989, <u>Foundations of Computer-Aided Design</u>, pp. 196-207, St. Paul, MN, USA, West Publishing Company.

9. Parkinson, B. W., Spilker, J. J., Eds., 1996, <u>Global Positioning System: Theory and Applications</u>, Vol. 1, pp. 547-68, Washington, DC, USA, American Institute of Aeronautics and Astronautics.

10. Townsend, B., Fenton, P., 1994, *A Practical Approach to the Reduction of Pseudorange Multipath Errors in a L1 GPS Receiver*, Proceedings of the 7^{th} International Technical Meeting of the Satellite Division of the Institute of Navigation, Part 1(of 2), Proceedings of ION GPS, Vol. 1, pp. 143-48.

11. Townsend, B., et al., 1995, *Performance Evaluation of the Multipath Estimating Delay Lock Loop*, NAVIGATION, Journal of Navigation, Vol. 42, No. 3, pp. 503-14.

12. Van Nee, D. J. R., 1992, *Multipath Effects on GPS Code Phase Measurements*, NAVIGATION, Journal of Navigation, Vol. 39, No. 2, pp. 177-90.