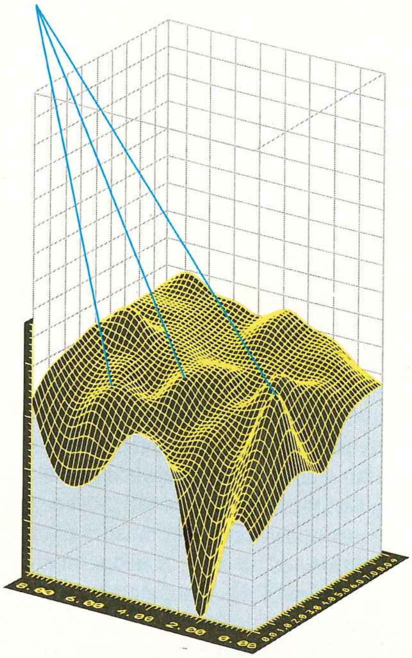

Conquering Multipath: The GPS Accuracy Battle

Lawrence R. Weill

California State University

In last month's column, we examined the error budget of stand-alone GPS positioning using the C/A-code. We saw that next to selective availability, the potentially largest contributor to the budget is multipath. This month, we will take a closer look at multipath and the techniques for mitigating its effects, including some recent innovative receiver designs. Our author is Lawrence R. Weill, a professor of applied mathematics at California State University, Fullerton, and a well-known GPS industry consultant who has operated his own consulting firm for 17 years. He is also one of the three technical founders of Magellan Systems Corporation. Dr. Weill has published numerous papers about signal-processing research for GPS, radar, sonar, optical sensor, and satellite communication systems. As an active GPS researcher, he has made substantial contributions to both the theoretical foundations and the practical aspects of GPS multipath mitigation.

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

The rapid evolution of GPS applications has produced a wide spectrum of performance requirements for GPS receivers, particularly in regard to their positioning accuracy. For users of low-cost, handheld receivers, 10–100-meter horizontal accuracy often suffices. At the other extreme, users of high-end survey-quality receivers can require centimeter-level accuracy in three dimensions, and in some cases, millimeter-level accuracy. The quest for ever higher accuracies has demanded a deeper understanding of GPS positioning error sources and how to reduce or eliminate them.

Differential GPS greatly reduces common-mode atmospheric, orbit, and satellite clock errors (including selective availability — SA). In addition, the last decade has seen much progress in reducing errors that occur within the GPS receiver itself. In fact, errors in receivers operating with ideal GPS signals have been reduced to near-theoretical limits.

As a result, designers now recognize that further error-reduction efforts must focus on multipath propagation, which produces errors that cannot be removed by differential operation. In recent years, certain user groups have recognized this as particularly crucial to their applications, including GPS surveying and the emerging Wide Area Augmentation System (WAAS), which is destined to serve the aviation community.

I will begin this article by defining the nature of the multipath problem and its implications for receiver performance. I will describe several spatially based methods for mitigating multipath errors. However, I will focus on real-time signal processing within the receiver because it offers the greatest promise and is today's most competitive research, development, and marketing arena. After surveying recent multipath-mitigation developments for C/A-code ranging, I will address an important question: How good can multipath mitigation get, and how close is current technology to those limits? I will also discuss the problem of reducing multipath errors in carrier-phase ranging, a matter of importance for high-accuracy applications. Before concluding, I will highlight the difficulties encountered when testing receiver performance.

THE MULTIPATH PROBLEM

All GPS receivers compute their position by determining the distance, or range, from the antenna to each of at least four (a minimum of three, if antenna height is known) GPS satellites, whose positions it computes from the broadcast ephemeris data. This set of distances (after they are corrected for satellite clock offsets, atmospheric propagation delay, and so forth) determines the point in space where the receiver is located. The receiver determines these distances by measuring each satellite signal's propagation time and then multiplying that number by the speed of light. The propagation time is determined by measuring the difference between transmission and reception times of the pseudorandom noise (PRN) code (for civilian use, the C/A-code) impressed on the signal. Survey-quality receivers also use propagation-time information conveyed by the signal carrier phase to achieve very accurate, although initially ambiguous, range measurements.

Because noise completely masks GPS signals received on earth, receivers use a method called correlation to process against the noise and accurately measure arrival time. Correlation will be discussed in more detail later in this article.

As long as each satellite's signal travels along a direct path straight to the receiver's antenna, the unit can determine the satellite range quite accurately. However, the ground and other objects easily reflect GPS signals, often resulting in one or more secondary paths, which are always longer than the direct path, as shown in Figure 1. These secondary-path signals, which are superimposed on the direct-path signals at the antenna, also have a longer propagation time and can significantly distort the signal waveform's amplitude and phase. Because a receiver without multipath protection requires an undistorted waveform, significant ranging errors can result.

In less expensive receivers that use only C/A-code ranging, a secondary-path signal simply produces a positioning error that can be tens of meters or more. High-end receivers, on the other hand, use both code- and carrier-phase ranging in differential GPS mode to achieve much higher accuracies. Initially ambiguous, these carrier phase–based ranges begin as a grid of possible values that must be searched and tested to find the correct one that will yield the receiver's true position. Accurate code ranging helps resolve the

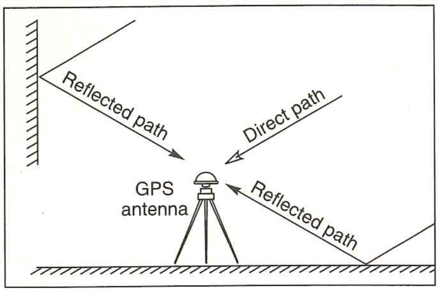

Figure 1. Multipath reflection geometry. GPS signals can be reflected from nearby structures or the ground. Ground bounce is a dominant scenario in practice.

ambiguity by limiting the search region, and accurate phase ranging thins out the number of eligible values within the region. Even small phase-measurement errors induced by multipath, therefore, can severely cripple this process, resulting in less-accurate positions.

Multipath propagation can be divided into two classes: static and dynamic. For a stationary receiver, the propagation geometry changes slowly, making the multipath parameters essentially constant for perhaps several minutes. But in mobile applications, a receiver can experience rapid fluctuations in fractions of a second.

However, the driving force for error reduction arises mostly from static applications, such as surveying, where greater demand for high-accuracy positioning exists. These applications often experience one dominant and stable secondary path — for example, from ground bounce or from a large structure in the antenna's vicinity. As a result, most research efforts have been within this context, even though some techniques are extendible to multiple secondary-path situations.

SPATIAL MITIGATION TECHNIQUES

Several multipath-reduction techniques attempt to take advantage of signal propagation geometry in one way or another. These methods include the use of special antennas (such as the choke-ring type), spatial processing with multiantenna arrays, antenna location strategies, and long-term signal observation to infer multipath parameters, facilitated by the changing reflection geometry. These methods try to reduce the strength of the secondary-path signals while preserving the direct-path signal — in other words, to isolate the direct-path signal.

Special Antennas. A simple form of multipath mitigation uses a metallic disk in the horizontal plane and centered at the GPS antenna's base. This technique's developers theorized that the disk, or extended groundplane, would shield the antenna from any signals arriving from below the antenna, such as those bouncing off the ground. However, this scheme did not perform as well as expected because of a quirky characteristic of electromagnetic waves. When a signal wavefront arrives at the disk's edge from below, it induces horizontally traveling surface waves on the disk's top side, which then travel to the antenna, thus compromising the disk's usefulness. Furthermore, not all multipath signals arrive from below the antenna, making this method ineffective in such cases.

To eliminate surface waves, the groundplane can be replaced with a choke ring — essentially a groundplane containing a series of concentric circular troughs one-quarter wavelength deep. These troughs act as transmission lines shorted at the far end and at their tops exhibit a very high impedance at the GPS signal frequency. Therefore, traveling surface waves cannot form, so the antenna gains a reasonable amount of protection from ground bounce and multipath signals arriving from near-horizontal directions.

The choke-ring antenna's disadvantage is that the circular troughs drive up its size, weight, and cost. Most importantly, the choke ring still cannot effectively mitigate multipath signals arriving from above the horizontal, as might be experienced from a reflection off a tall building. Nonetheless, such antennas have proven themselves in applications where ground bounce provides the dominant multipath source, particularly in surveying.

A GPS antenna designed for the right-hand circularly polarized (RHCP) signals transmitted by GPS satellites should provide some degree of immunity to multipath signals arising from reflection. In theory, a RHCP signal becomes left-hand circularly polarized (LHCP) upon reflection from an ideal conducting surface, and the RHCP antenna is not nearly as responsive to the LHCP signal.

Multiantenna Spatial Processing. Users can also reduce multipath effects by deploying multiple antennas that simultaneously receive the GPS signal at different points in space. Because the multipath geometry varies at different spatial locations, the multipath-corrupted GPS signal will generally have different characteristics at each antenna. Users can employ a form of signal processing called *spatial processing* to exploit these differences and isolate the desired direct-path signal. In some cases, multiple-antenna use can be thought of as forming a directional antenna pattern responsive to the direct-path signal but not to multipath signals arriving from other directions.

Antenna Location Strategy. Users can often greatly reduce multipath effects by placing the antenna where it is less likely to receive reflected signals. For example, to position a point near a potentially reflective object, one could determine the position of a nearby point "in the clear" and calculate the desired position based on offsets obtained using terrestrial distance and/or angle-measuring techniques. Another of my favorite examples eliminates ground bounce reflection in GPS surveying by placing the receiver antenna on the ground instead of a tripod. Clearly, antenna relocation may be impractical in some cases, but it can be effective when feasible.

Long-term Signal Observation. If a receiver observes a signal for a long time, from sizeable fractions of an hour to hours, one can take advantage of the changing geometry of the secondary-path reflections, caused by the GPS satellites' angular motion across the sky. This motion causes the relative delays between the primary and secondary propagation paths to change, causing measurable variations in the received signal.

Some approaches remove the secondary-path components by identifying them through the variations in signal level or signal-to-noise ratio caused by alternate phase reinforcement and cancellation. Because they require long signal-observation times, these techniques are specialized and would probably be impractical for a majority of applications.

One technique that can be used to accurately characterize multipath at a fixed site, such as at a differential GPS base station, is to observe the same satellites from one day to the next, looking for patterns in pseudorange or phase measurements that are advanced by about four minutes per day because of the satellites' nominal half-a-sidereal-day period.

RECEIVER PROCESSING METHODS

By far the most promising methods for reducing multipath effects use real-time signal processing within the receiver. This new and exciting cutting-edge technology involves a flurry of research and development by major receiver manufacturers that often tout the results of their approaches without explicitly revealing their secrets.

The work's proprietary nature poses a real challenge when reporting on this field's latest developments, and the level of information companies provide about their technology varies greatly. Some companies permit technical paper presentations, but these often omit proprietary information and reveal only results. Others simply give a special name to their technology and issue vague statements about how well it can reduce multipath.

Despite this complicated state of affairs, enough information about multipath processing exists to gain insight into its recent evolution. Most published work has dealt with range measurement using the C/A-code, which is what I will treat in the remainder of this section. However, mitigation of multipath effects on the more precise carrier-phase range measurements remains important and will be briefly discussed later.

Standard Range Measurements. As a foundation for investigating various real-time signal processing techniques that mitigate multipath, let's first explore how multipath causes errors in receivers that use standard pseudoranging methods — that is, in receivers that are not specially designed to handle multipath signals. For this purpose, I will make some simplifications that in no way obscure the fundamentals involved. We will assume that the receiver processes only the C/A-code after the received signal has been frequency shifted to baseband (nominally zero frequency). When no multipath is present, the receiver waveform at reception time, *t*, can be mathematically represented in complex-notation (phasor) form by

$$r(t) = ae^{j\phi}c(t - \tau) + n(t) \qquad [1]$$

where $c(t)$ is the normalized, undelayed C/A-code waveform as transmitted, τ is the signal propagation delay, a is the signal amplitude, ϕ is the carrier phase, and $n(t)$ is receiver thermal noise. (For those who may have forgotten their high-school complex number theory, e is the base of the natural logarithms and j is the square root of minus one.)

For simplicity, we also assume that the signal has been compensated for Doppler shift and that the 50 bit-per-second data modulation on the signal has been removed by standard techniques. Pseudoranging aims to accurately estimate the propagation delay τ, which can then be converted into the satellite range. In principle, the receiver makes this estimate by generating a replica $c_r(t)$ of the transmitted C/A-code and aligning the replica in time so as to match the arriving code waveform as closely as possible. The degree of alignment is measured by the magnitude of the cross-correlation function of the incoming signal and the receiver-generated replica.

The cross-correlation function can be expressed mathematically as

$$R(\hat{\tau}) = \int_{T_1}^{T_2} r(t)c_r(t - \hat{\tau})dt \qquad [2]$$

where $R(\hat{\tau})$ is the cross-correlation function, $\hat{\tau}$ is the time shift of the receiver code replica $c_r(t)$, and $r(t)$ is the noisy received waveform given by equation [1]. The integration over the time interval from T_1 to T_2 is chosen to provide a large amount of processing gain that greatly reduces error caused by thermal noise in the receiver. In most receivers, the effective integration time is roughly the reciprocal of the code-tracking loop bandwidth. Typical GPS signals require an integration time of roughly 1 second (loop bandwidth of 1 Hz) to get C/A-code range estimates that are precise to within 1–2 decimeters. The estimate of the actual signal propagation delay τ is computed by varying $\hat{\tau}$ until the magnitude of $R(\hat{\tau})$ is maximum, which happens when the received and replica waveforms become aligned in time, except for an error caused by noise.

In most receivers the correlation function value is represented digitally with enough

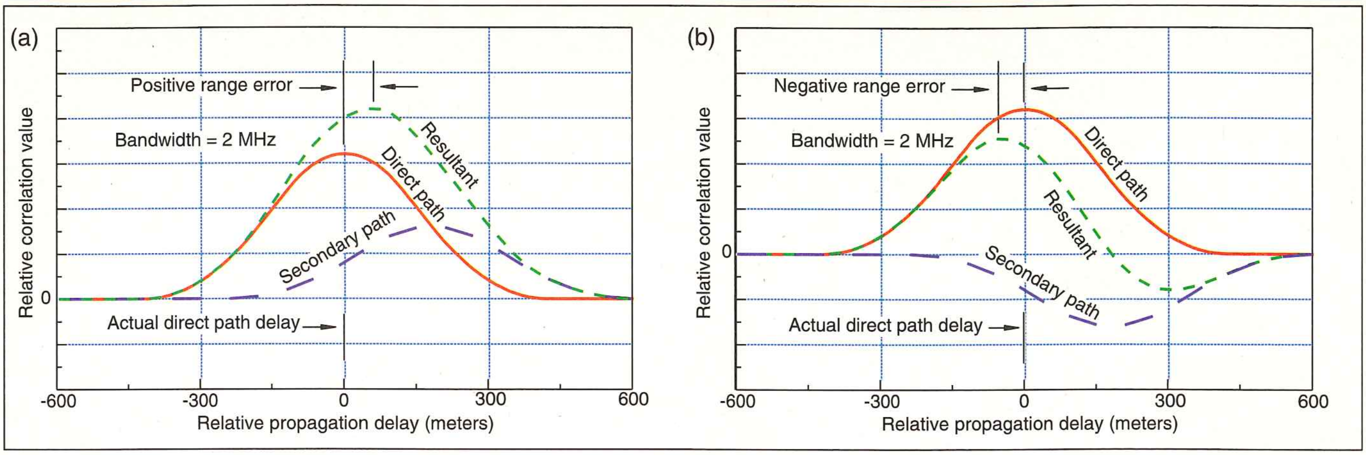

Figure 2. (a) Positive range error caused by an in-phase secondary signal path. The positive slope of the secondary path cross-correlation function shifts the peak of the direct path cross-correlation function to the right, as shown by the resultant curve. **(b)** Negative range error occurs when the secondary path is out of phase with the direct path.

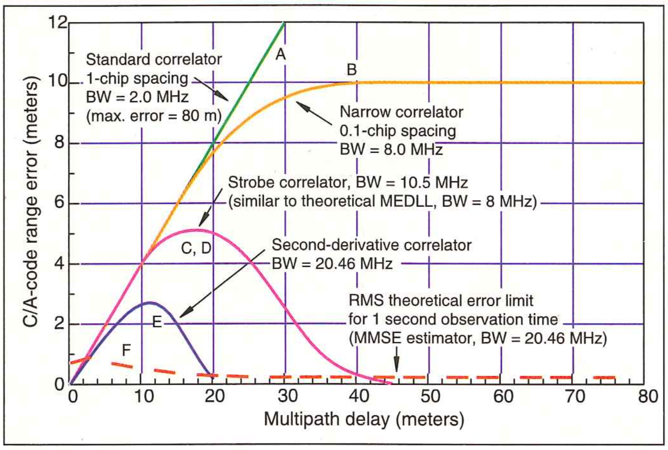

Figure 3. Typical code ranging errors for several multipath mitigation methods. Curves are for a secondary path relative phase of zero degrees and show bias (in other words, systematic) errors, except for the theoretical error bound, which is a root-mean-square (rms) error. In all cases the secondary path amplitude is one-half that of the direct path.

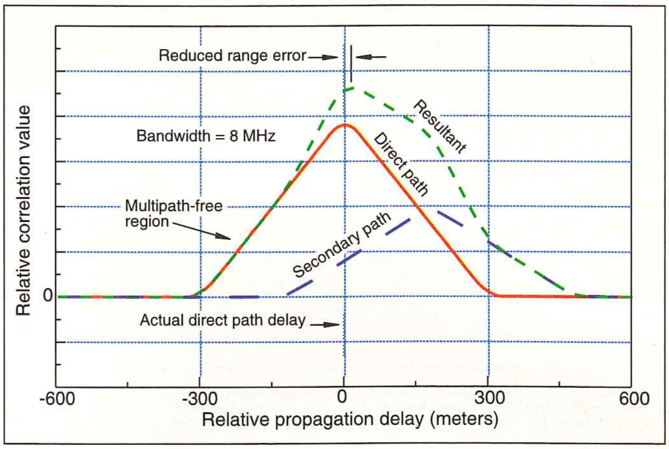

Figure 4. Increasing the precorrelation receiver bandwidth sharpens the peak of the direct-path cross-correlation function, thereby reducing the error caused by a secondary signal. Note that the leading portion of the resultant cross-correlation function is uncontaminated by multipath.

bits to allow sufficiently accurate representation of its value. Generally speaking, the bit quantization size needs to be significantly smaller than the standard deviation of the noise on the correlation function to accurately locate the peak of the correlation function for ranging purposes.

A typical cross-correlation function without multipath, for receivers having a 2-MHz precorrelation bandwidth (the C/A-code, with a chipping rate of approximately 1 megabit per second, has a principal spectral lobe that occupies a bandwidth of about 2 MHz) is shown by the solid-line curves of Figures 2a and 2b. (The plots ignore the effect of noise, which would add small random variations to the curves.) For greatest pseudoranging accuracy, the cross-correlation function's peak should be as sharp as possible so that the noise will have as little effect as possible on determining its location. The peak's sharpness (given by the magnitude of its second derivative — the rate of change of its slope) can be shown to increase with the signal's bandwidth before correlation. Surprisingly, this fact was just recently explicitly exploited in GPS receiver design (see the following discussion of narrow-correlator technology).

If multipath is present with a single secondary path, the waveform of equation [1] changes to

$$r(t) = ae^{j\phi_1}c(t - \tau_1) + be^{j\phi_2} c(t - \tau_2) + n(t). \quad [3]$$

In this signal, the direct and secondary paths have respective propagation delays τ_1 and τ_2, amplitudes a and b, and carrier phases ϕ_1 and ϕ_2. In a receiver not designed to expressly handle multipath, the resulting cross-correlation function will now have two superimposed components, one from the direct path and one from the secondary path, as illustrated in Figures 2a and 2b. The result is no longer a nice symmetric triangle, but a distorted curve.

Furthermore, the location of this function's peak, which the receiver unwittingly interprets as the direct-path range, is displaced from its correct position because of the secondary-path component's slope. The nature of the resulting distortion depends on the secondary-path component's strength, location, and relative phase. Figures 2a and 2b respectively illustrate the effects of secondary-path in-phase and out-of-phase conditions in which the secondary path has one-half the direct path's amplitude (larger-amplitude secondary-path signals would result in larger multipath errors).

In 1973, L.L. Hagerman of the Aerospace Corporation published a report in which he discussed the dependence of the resulting range error in vintage receivers employing standard code-tracking techniques. Curve A of Figure 3 illustrates this dependence for the case where the direct and secondary paths are in phase reinforcement. The error's magnitude is essentially zero at zero path separation, increases to a maximum value of 80 meters or so at about 250 meters of path separation, and then declines to zero when the path separation is enough to give essentially complete C/A-code decorrelation.

A Correlation Function's Leading Edge. Several researchers have proposed a multipath reduction technique that takes advantage of the fact that direct-path signals always arrive before those from longer secondary paths. This means multipath does not contaminate the leading portion of the cross-correlation function, as can be seen by looking at the leftmost part of the resultant cross-correlation curve in Figure 4.

Therefore, if one could measure the location of just the leading part of the cross-correlation function, it appears that the signal delay could be determined and multipath would cause no errors at all! However, this seemingly happy state of affairs is somewhat illusory. With a small direct-to-secondary path separation, the cross-correlation function's uncontaminated portion is just a minuscule piece at the extreme left, where the curve just begins to rise. In this region, not only is the signal-to-noise ratio relatively poor, but the curve's slope is also relatively small, which makes it difficult to accurately estimate the delay associated with this multipath-free region.

For these reasons, this approach best suits situations with a moderate to large direct-to-secondary path separation. However, even in these cases, an additional problem must be addressed. Because the receiver must estimate the delay of the leading part (and not the peak) of the cross-correlation function, the receiver must employ some method to make the delay measurement insensitive to the slope of the correlation function's leading edge, which can vary with signal strength and other factors. Such a problem does not occur in multipath-mitigation methods that estimate only the location of the cross-correlation function's peak.

Narrow-Correlator Technology (1990–93). The first significant means to reduce GPS multipath effects by receiver processing was probably discovered when *narrow-correlator* technology made its debut in the early 1990s. Until that time, receiver designers matched the front-end bandwidth of most receivers to the spread-spectrum GPS signal's bandwidth (about 2 MHz, as we mentioned earlier). However, at this bandwidth, the pseudorange cross-correlation function's peak is very rounded. The heavy rounding allows a secondary-path component to cause a large shift in the resulting function's peak away from the position that correctly represents direct-path range, as previously illustrated in Figures 2a and 2b.

An important paper, published in 1992 by A.J. Van Dierendonck and others, clarified matters by showing that for best ranging accuracy, one should use a wide precorrelation bandwidth rather than a narrow one, a notion that defied the intuition of some designers. However, we now understand that the sharpening of the correlation function's peak afforded by a wider bandwidth (as wide as 20 MHz) will greatly reduce the effect of noise in determining its location.

Algorithms that locate and track the sharpened peak are called *narrow correlators* because they employ two code replicas with closely spaced time delays to straddle the narrowed-correlation function peak. Because of inherent noise cancellation properties, the ranging accuracy of narrow correlators turned out to be quite impressive, making errors on the order of only 20–30 centimeters theoretically possible, with healthy GPS signals and no multipath.

Narrow-correlator development appeared to be motivated by the quest to reduce range-estimation errors caused by receiver thermal noise rather than those caused by multipath. However, it was soon discovered that the narrow correlator offers an additional benefit — significantly better multipath performance than that obtained with standard correlation. The reason is not hard to explain: the location of a sharp peak in the direct-path correlation component is less easily shifted by the presence of a secondary component, as illustrated in Figure 4.

Curve B of Figure 3 shows typical multipath-induced error for the narrow correlator as a function of path separation when the two paths are in phase reinforcement. This might be considered the first step in receiver multipath processing technology, in which the 80-meter worst-case error of curve A has been reduced to about 10 meters.

Correlation-Function Shapes (1994–95). Observation of the correlation functions of Figures 2 and 4 gives rise to an idea that has been exploited by at least one major GPS receiver manufacturer: the cross-correlation function's shape depends on the multipath parameters. Thus, it might be possible to determine these parameters from the specific, observed shape, and in so doing obtain an accurate estimate of the direct-path propagation delay parameter. The idea has merit but requires many correlations with different values of signal replica delay $\hat{\tau}$ to obtain a sampled version of the function shape. Various methods can be used to map this shape into the corresponding multipath parameters.

A straightforward approach is to store or compute a number of amplitude-normalized correlation functions corresponding to different multipath parameter combinations (path amplitudes, delays, and phases) and find the best match to the normalized correlation function actually observed. Another technique, called the early-late slope method, calculates a pseudorange correction by measuring the differing slopes caused by multipath, on either side of the correlation function peak.

Curve C of Figure 3 is representative of the results obtained by one manufacturer using its own version of shape discrimination. Called MEDLL (for multipath estimating delay-lock loop), it represents a significant improvement over the narrow-correlator technology alone, not only because the worst-case ranging error is generally smaller (about 6 meters versus 10), but because significant ranging error occurs over a much smaller interval of path separations. The MEDLL algorithms employed are based on the work of Richard van Nee of the Delft University of Technology, who used the statistical method of maximum likelihood to estimate the multipath parameters. For more details on MEDLL, consult the "Further Reading" sidebar.

The Strobe Correlator. At the Institute of Navigation Satellite Division's 9th International Technical Meeting, ION GPS-96, a leading receiver manufacturer presented a paper describing the results of a proprietary multipath-mitigation technique called the *strobe correlator*. Curve D of Figure 3 depicts this technique's performance, which is similar to the theoretical results for MEDLL. The strobe correlator shares with MEDLL the advantage that, unlike the narrow correlator, it almost completely eliminates multipath errors for delays exceeding about 40 meters. To my knowledge, no details have been published on exactly how the strobe correlator achieves the results shown. However, statements in the referenced paper suggest that an unusual correlator reference waveform is used.

Modified Correlator Reference Waveforms. Another approach to multipath mitigation, so new that it is apparently not well-known by the GPS community, alters the waveshape of the correlator reference PRN code to provide a cross-correlation function that has inherent resistance to errors caused by multipath. Here I shall discuss one example of this generic method.

Referring again to Figure 4, we see that at the peak of the direct-path cross-correlation function (solid curve), the function's second derivative is maximum in a negative direction (the second derivative is closely related to the amount of curvature of the graph). This leads to the idea of measuring signal delay by forming the negative of the cross-correlation function's second derivative and finding the location of *its* peak. We can form the second derivative of the cross-correlation function in an optimal receiver by replacing the correlator reference code waveform (which is a band-limited version of an ideal waveform) with its second derivative. When this is done, equation [2] for the cross-correlation function changes to

$$R''(\hat{\tau}) = \int_{T_1}^{T_2} r(t)c_r''(t-\tau)dt \qquad [4]$$

where $R''(\hat{\tau})$ is the second derivative of the cross-correlation function and $c_r''(t)$ is the second derivative of the reference code.

The advantage of finding the peak of $-R''(\hat{\tau})$ instead of $R(\hat{\tau})$ is that the sloping portion of the secondary-path cross-correlation component, as shown in Figure 4, has a second derivative equal to zero, which means that it can no longer shift the direct-path component's peak. Therefore, errors caused by multipath can be greatly reduced for a wide range of path separations. Curve E of Figure 3 shows a typical performance curve for this approach. This demonstrates that this method remains most effective for path separations greater than about 20 meters.

HOW GOOD CAN IT GET?

An important consideration in the continuing quest for better and better multipath performance is the state of the art in relation to the best performance that is theoretically possible. Discovering such performance limits would provide a valuable benchmark in receiver design.

For example, if we knew a particular algorithm's performance was close to a theoretical limit, futile and costly attempts to significantly improve performance could be avoided. More importantly, finding the theoretically optimum performance invariably leads to a method for achieving it and could provide the basis for an "unbeatable" algorithm. Of course, the problem of implementing it in a feasible manner might remain.

Surprising and interesting subtleties arise in attempts to determine the best possible multipath mitigator. For example, it can be proven that no algorithm exists that is best under all multipath conditions! However, it turns out that one estimator for mitigating multipath can be claimed as optimal in a certain sense. Called the *minimum-mean-square error* (MMSE) *estimator*, this method relies on a branch of statistics known as Bayesian estimation theory.

Further Reading

For the original paper on narrow-correlator theory, see

- "Theory and Performance of Narrow Correlator Spacing in a GPS Receiver," by A.J. Van Dierendonck, P. Fenton, and T. Ford, published in the *Proceedings of The Institute of Navigation National Technical Meeting*, held in San Diego, California, January 27–29, 1992, pp. 115–124.

For the development of the early-late slope (ELS) and multipath estimating delay-lock loop (MEDLL) techniques, refer to

- "A Practical Approach to the Reduction of Pseudorange Multipath Errors in a L1 GPS Receiver," by B. Townsend and P. Fenton, published in the *Proceedings of ION GPS-94, the 7th International Technical Meeting of the Satellite Division of The Institute of Navigation*, held in Salt Lake City, Utah, September 20–23, 1994, pp. 143–148.
- "Performance Evaluation of the Multipath Estimating Delay-Lock Loop," by B. Townsend, D.J.R. van Nee, P. Fenton, and K. Van Dierendonck, published in the *Proceedings of The Institute of Navigation National Technical Meeting*, held in Anaheim, California, January 18–20, 1995, pp. 277–283.
- *Multipath and Multi-Transmitter Interference in Spread-Spectrum Communication and Navigation Systems*, by R.D.J. van Nee, published by Delft University Press, Delft, The Netherlands, 1995.

A description of strobe correlator results can be found in

- "Strobe & Edge Correlator Multipath Mitigation for Code," by L. Garin, F. van Diggelen, and J. Rousseau, published in the *Proceedings of ION GPS-96, the 9th International Technical Meeting of the Satellite Division of The Institute of Navigation*, held in Kansas City, Missouri, September 17–20, 1996, pp. 657–664.

Details about the use of special pseudorandom-noise code reference waveforms to achieve multipath mitigation are presented in

- "GPS Multipath Mitigation by Means of Correlator Reference Waveform Design," by L. Weill, to be published in the *Proceedings of the National Technical Meeting of The Institute of Navigation*, held in Santa Monica, California, January 14–16, 1997.

Theoretical performance limits for multipath error reduction are treated in

- "Achieving Theoretical Accuracy Limits for Pseudoranging in the Presence of Multipath," by L. Weill, published in the *Proceedings of ION GPS-95, the 8th International Technical Meeting of the Satellite Division of The Institute of Navigation*, held in Palm Springs, California, September 12–15, 1995, pp. 1521–1530.

Carrier-phase multipath mitigation results can be found in

- "L1 Carrier Phase Multipath Error Reduction Using MEDLL Technology," by B. Townsend, P. Fenton, K. Van Dierendonck, and R. van Nee, published in the *Proceedings of ION GPS-95, the 8th International Technical Meeting of the Satellite Division of The Institute of Navigation*, held in Palm Springs, California, September 12–15, 1995, pp. 1539–1544.

MMSE treats the multipath parameters as random variables and uses the signal observed by the receiver to construct a conditional probability density for the parameter values. The optimality property for this estimator is that no other estimator has a uniformly smaller root-mean-square (rms) error. In other words, if some other direct-path range estimator performs better than the MMSE estimator under certain multipath conditions, then that estimator must perform more poorly than the MMSE estimator under some *other* multipath conditions.

Curve F of Figure 3 shows the code ranging performance of a two-path MMSE estimator incorporating all significant multipath parameters, where a uniform *a priori* parameter probability density has been assumed and the signal observed for one second. Under most conditions, this estimator performs better than the previously described estimators, with about 0.9-meter rms worst-case error. This suggests that we are not yet at the end of the technology improvement road.

Why haven't designers incorporated the MMSE estimator into a GPS receiver? Perhaps they have, and we haven't been told. However, that seems unlikely, as its direct implementation is *extremely* computation intensive, and uncommon cleverness would be required to make it feasible. On the positive side, the MMSE estimator has a very desirable property: one can decrease range error simply by observing the signal for a longer time period. It remains unclear whether other recently developed methods offer this advantage.

CARRIER-PHASE RANGING

The presence of multipath also has deleterious effects on carrier-phase range measurements, which limits the performance of high-end GPS receivers used for surveying and other precision applications, particularly with regard to carrier-phase ambiguity resolution. Despite its importance, it is remarkable that relatively little published work shows clearly defined, in-receiver mitigated carrier-phase error curves. (See the "Further Reading" sidebar for a published paper that includes carrier-phase error results.)

Although some receiver manufacturers have made rather vague performance claims, one gets the impression that mitigation of multipath carrier-phase error presents a substantially more difficult problem than pseudorange-error mitigation. My own research indicates that the most difficult situation for reducing carrier-phase error occurs at small path separations, and essentially no mitigation is possible when the separation is about a meter or less. We hope to soon see more definitive results in this important area.

RECEIVER TESTING

Conducting meaningful tests of receiver multipath-mitigation performance, on either an absolute or comparative basis, is no easy matter. Two conflicting goals often exist. On one hand, the testing should be under strictly controlled conditions, so that the signal levels and true multipath parameters are precisely known; otherwise one cannot tie the measured performance to the multipath conditions that actually exist. Generally this requires precision signal simulators and other hardware to generate accurately characterized multipath signals.

On the other hand, receiver end users place more credence in how much improve-

Continued on page **66**

NOTES

Motorola Adds OEM Chipset

Motorola has introduced a chipset version of its eight-channel Oncore GPS technology, targeting original equipment manufacturers (OEMs) and systems integrators.

The chipset, expected to be available in May, includes a radio-frequency integrated circuit (RFIC) and digital correlator IC optimized for Motorola's 68331 32-bit microcontroller, along with Oncore software. Motorola will offer the software in two versions: GT Oncore optimized for land-vehicle positioning and tracking applications, and UT Oncore, for precise-timing and frequency-stabilization applications.

The original Basic Oncore, a six-channel module introduced in 1992, marked the diversified $26-billion electronic company's focus on OEM products and markets in building a GPS business.

Trimble Inks AVL Agreements

Trimble Navigation has signed an agreement with Xanavi Informatics Corporation to provide access to Trimble's eight-channel Sierra GPS chipset for use in the Japanese company's in-vehicle navigation systems. The technology will be licensed for an undisclosed royalty.

Xanavi, established in 1991 by Hitachi and Nissan Motor Company, operates as an integrated developer, producer, and marketer of automotive information and communications systems, including the GPS-enhanced Birdview car navigation and information system. Trimble previously has supplied GPS boards and chipsets to Xanavi.

Trimble also recently announced its receipt of a $1.5-million contract to provide a GPS-based automatic vehicle location system (AVL) services to the city of Houston, Texas. The contract calls for Trimble to equip more than 200 vehicles in the city's fire and emergency medical service departments.

SiRF Moves, Raises Funds

SiRF Technology, Inc. has closed an $8-million round of financing to provide working capital and build inventories as it ramps up production for its SiRFstar GPS chipsets. The start-up company has also relocated its headquarters from Sunnyvale, California, to Santa Clara, California, and expanded its research and development facilities in Los Angeles. The new investors include three Taiwan-based financial companies: Fortune Venture Capital Company, InveStar Capital Inc., and Pacific Venture Partners.

The company also received investments from Singapore-based ECICS Ventures Limited and a number of private and current investors. Existing investors include Ayala Corporation; FPHC International; Mitsui & Company, Yamaha Corporation; and the Walden Group.

SiRF Technology can now be contacted at 3970 Freedom Circle, Santa Clara, CA 95054, USA, (408) 980-4700, fax (408) 908-4705. Further information about the company can also be found at its new Web site, <http://www.sirf.com.>

Innovation at Work

GPS World's popular Applications Contest, an annual feature in the August *Showcase* issue, is accepting entries for the 1997 competition.

The contest provides readers a series of "snapshots" of innovative GPS applications submitted by users. Winning entries receive prizes donated by GPS manufacturers. Although entries are not accepted from employees of GPS companies or commercial system integrators, the magazine does encourage these organizations to notify their customers about the contest and support their efforts to develop contest entries.

Participants need to send a brief, typewritten description of their work (no more than 300 words) along with slides or color photographs illustrating the application; information on the equipment used, and entrants' affiliation, location, and full contact information. An independent panel of judges will evaluate the edited submissions based on three criteria: innovation, technology, and practical results. Innovation will be the main criterion. All acceptable entries will be published in the August 1997 *GPS World Showcase.*

Project descriptions should be sent before May 30 to: Ling Chan, managing editor, c/o GPS World, 859 Willamette Street, Eugene, OR 97401-6806, USA; (541) 984-5247; fax (541) 344-3514; or e-mail <editorial-gps@gpsworld.com>.

Around the Industry . . .

NovAtel Inc. (Calgary, Alberta, Canada) has signed an agreement with the **Norman Wade Company Limited** (Calgary) to distribute NovAtel products in Canada. Norman Wade maintains 18 locations throughout Canada, supplying products for the survey, engineering, architectural, and related industries. ■

Continued from page **65**

ment is noted in the field. However, this is almost impossible to measure accurately, because one does not know the amount and character of the multipath and can encounter great difficulty in isolating errors caused by multipath from those of other sources. The question also arises as to whether the actual multipath conditions fall within the assumptions used in the design of the multipath-mitigation algorithm. For example, if three propagation paths are actually present, the performance of an algorithm designed for two paths might be questionable.

In spite of these assessment difficulties, one will likely be able to note, after long periods of use under various conditions, the difference in performance between receivers that employ and do not employ multipath mitigation.

CONCLUDING REMARKS

Without doubt, the progress in GPS multipath mitigation has been significant and will continue. In view of the disparity between what is theoretically achievable and what appears to be the state of the art, there seems to be room for more rigorous approaches. Definitive results in the mitigation of carrier-phase errors are also less numerous, as compared with more extensively published results related to code-ranging errors. Despite the fact that using the full, GPS C/A-code bandwidth (about 20 MHz, encompassing many spectrum sidelobes) provides the best ranging accuracy and multipath mitigation, most currently sold receivers use a narrower bandwidth. This will undoubtedly change in the near future.

One final comment: The reluctance of GPS receiver manufacturers to reveal their proprietary methods of multipath error reduction is understandable. But it is hoped that more information on the subject will become available not just to sophisticated GPS users, but to the scientific community at large, because its applications extend beyond GPS. Communications, sonar, and radar systems, for example, could also benefit from what has already been learned and what will be learned in the future. ■

MANUFACTURERS

The narrow-correlator and MEDLL technologies were developed or adapted by **NovAtel Inc.** (Calgary, Alberta, Canada). The strobe correlator was developed by **Ashtech, Inc.** (Sunnyvale, California).