

Collective Detection

Enhancing GNSS Receiver Sensitivity by Combining Signals from Multiple Satellites

Penina Axelrad, James Donna, Megan Mitchell, and Shan Mohiuddin

ALTHOUGH I HAVE MANAGED THE INNOVATION COLUMN continuously since *GPS World's* first issue, it wasn't until the second issue that I authored a column article. That article, co-written with Alfred Kleusberg, was titled "The Limitations of GPS." It discussed some of the then-current problems of GPS, including poor signal reception, loss of signal integrity, and limited positioning accuracy.

In the ensuing 20 years, both signal integrity and positioning accuracy have improved significantly. Advances in the GPS control segment's capabilities to continuously monitor and assess signal performance, together with receiver-autonomous integrity monitoring and integrity enhancement provided by augmentation systems, have reduced worries about loss of signal integrity.

The removal of Selective Availability and use of error corrections provided by augmentation systems, among other approaches, have improved positioning accuracy.

But the problem of poor reception due to weak signals is still with us. In that March/April 1990 article, we wrote "[GPS] signals propagate from the satellites to the receiver antenna along the line of sight and cannot penetrate water, soil, walls, or other obstacles very well. . . . In surface navigation and positioning applications, the signal can be obstructed by trees, buildings, and bridges. . . . [In] the inner city streets of urban areas lined with skyscrapers, the 'visibility' of the GPS satellites is very limited. In such areas, the signals can be obstructed for extended periods of time or even [be] continuously unavailable."

Poor signal reception in other than open-sky environments is still a problem with conventional GPS receivers. However, extending signal integration times and using assisted-GPS techniques can give GPS some degree of capability to operate indoors and in other restricted environments, albeit typically with reduced positioning accuracy. An antenna with sufficient gain is needed and capable systems are available on the market. The pilot channels of modernized GNSS signals will also benefit signal acquisition and tracking in challenging environments.

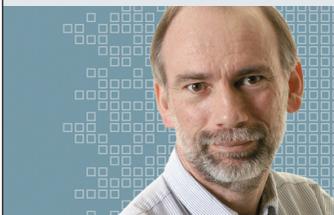
In this month's column, we look at a completely different approach to enhancing signal sensitivity. Rather than requiring each satellite's signal to be acquired and tracked before it can be used in the navigation solution, the new approach — dubbed "collective detection" — combines the received signal power from multiple satellites in a direct-to-navigation-solution procedure. Besides providing a quick coarse position solution with weak signals, this approach can be used to monitor the signal environment, aid deeply-coupled GPS/inertial navigation, and assist with terrain and feature recognition.

"Innovation" features discussions about advances in GPS technology, its applications, and the fundamentals of GPS positioning. The column is coordinated by Richard Langley, Department of Geodesy and Geomatics Engineering, University of New Brunswick. To contact him with topic ideas, see "Contributing Editors," page 6.

Growing interest in navigating indoors and in challenging urban environments is motivating research on techniques for weak GPS signal acquisition and tracking. The standard approach to increasing acquisition and tracking sensitivity is to lengthen the coherent integration times, which can be accomplished by using the pilot channels in the modernized GPS signals or by using assisted GPS (A-GPS) techniques. These techniques operate in the traditional framework of independent signal detection, which requires a weak signal to be acquired and tracked before it is useful for navigation. This article explores a complementary, but fundamentally different, approach that enhances signal sensitivity by combining the received power from multiple GPS satellites in a direct-to-navigation-solution algorithm. As will be discussed in the following sections, this collective detection approach has the advantage of incorporating into the navigation solution information from signals that are too weak to be acquired and tracked, and it does so with a modest amount of computation and with no required hardware changes. This technology is appropriate for any application that requires a navigation solution in a signal environment that challenges traditional acquisition techniques. Collective detection could be used to monitor the signal environment, aid deeply coupled GPS/INS during long outages, and help initiate landmark recognition in an urban environment. These examples are explained further in a subsequent section. In order to understand how the collective detection algorithm works, it is instructive to first consider the traditional approach to acquisition and tracking.

Acquisition Theory and Methods

In a typical stand-alone receiver, the acquisition algorithm assesses the signal's correlation power in discrete bins on a

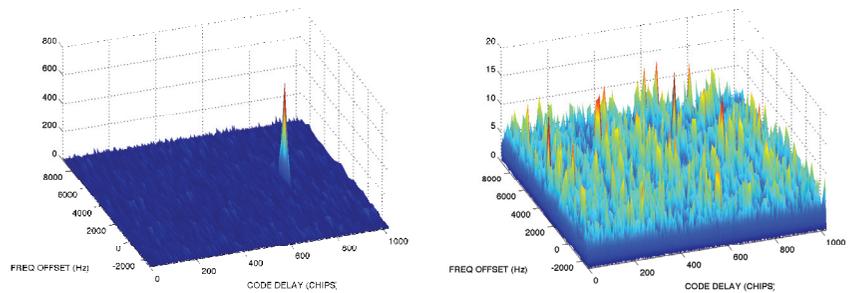


INNOVATION INSIGHTS
with Richard Langley

Poor signal reception in difficult environments is still a problem.

grid of code delay and Doppler frequency (shift). The correlation calculations take the sampled signal from the receiver's RF front end, mix it with a family of receiver-generated replica signals that span the grid, and sum that product to produce in-phase (I) and quadrature (Q) correlation output. The correlation power is the sum of the I and Q components, $I^2 + Q^2$. Plotting the power as a function of delay and frequency shift produces a correlogram, as shown in **FIGURE 1**. It should be noted that both correlation power and its square root, the correlation amplitude, are found in the GPS literature. For clarity, we will always use the correlation power to describe signal and noise values.

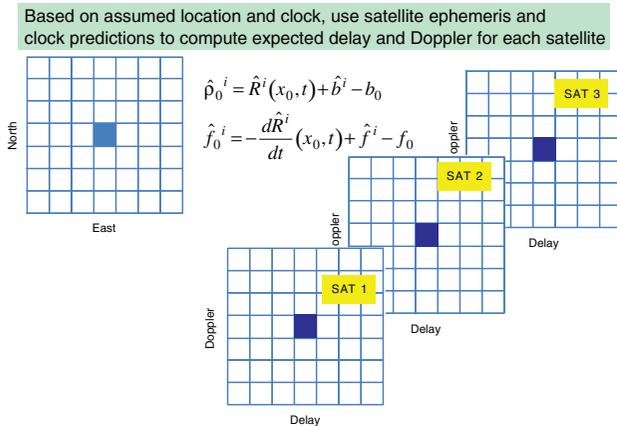
If a sufficiently powerful signal is present, a distinct peak appears in the correlogram bin that corresponds to the GPS signal's code delay and Doppler frequency. If the peak power exceeds a predefined threshold based on the integration times



▲ **FIGURE 1** Correlograms for a strong GPS signal (left) and a weak GPS signal (right).

and the expected carrier-to-noise spectral density, the signal is detected. The code delay and Doppler frequency for the peak are then passed to the tracking loops, which produce more precise measurements of delay — pseudoranges — from which the receiver's navigation solution is calculated. When the satellite signal is attenuated, however, perhaps due to foliage or building materials, the correlation peak cannot be distinguished and the conventional approach to acquisition fails.

The sensitivity of traditional tracking algorithms is similarly limited by the restrictive practice of treating each signal independently. More advanced tracking algorithms, such as vector delay lock loops or deeply integrated filters, couple the receiver's tracking algorithms and its navigation solution in order to take advantage of the measurement redundancy and to leverage information gained from tracking strong signals to track weak signals. The combined satellite detection approach presented in



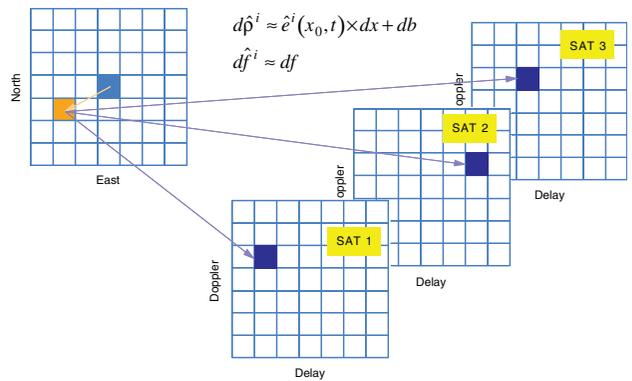
▲ FIGURE 2 Centering the various search spaces on the *a priori* values.

this article extends the concept of coupling to acquisition by combining the detection and navigation algorithms into one step.

Collective Detection

In the collective detection algorithm, a receiver position and clock offset grid is mapped to the individual GPS signal correlations, and the combined correlation power is evaluated on that grid instead of on the conventional independent code delay and Doppler frequency grids. The assessment of the correlation power on the position and clock offset grid leads directly to the navigation solution. The mapping, which is key to the approach, requires the receiver to have reasonably good *a priori* knowledge of its position, velocity, and clock offset; the GPS ephemerides; and, if necessary, a simplified ionosphere model. Given this knowledge, the algorithm defines the position and clock offset search grid centered on the assumed receiver state and generates predicted ranges and Doppler frequencies for each GPS signal, as illustrated in FIGURE 2. The mapping then relates each one of the position and clock offset grid points to a specific code delay and Doppler frequency for each GPS satellite, as illustrated in FIGURE 3. Aggregating the multiple delay/Doppler search spaces onto a single position/clock offset search space through the mapping allows the navigation algorithm to consider the total correlation power of all the signals simultaneously. The correlation power is summed over all the GPS satellites at each position/clock-offset grid point to create a position domain correlogram. The best position and clock-offset estimates are taken as the grid point that has the highest combined correlation power. This approach has the advantage of incorporating into the position/clock-offset estimate information contained in weak signals that may be undetectable individually using traditional acquisition/tracking techniques.

It should be noted that a reasonable *a priori* receiver state estimate restricts the size of the position and clock-offset grid such that a linear mapping, based on the standard measurement sensitivity matrix used in GPS positioning, from the individual signal correlations, is reasonable. Also, rather than attempt to align the satellite correlations precisely enough to perform coherent sums, noncoherent sums of the individual satellite correlations are used. This



▲ FIGURE 3 Map points from the position and clock-offset search space to code-delay and Doppler-frequency search space for several satellites. For short acquisitions with a stationary user and search space that is not too large, the Doppler-frequency offset from the *a priori* values is constant for each satellite.

seems reasonable, given the uncertainties in ranging biases between satellites, differences and variability of the signal paths through the ionosphere and neutral atmosphere, and the large number of phases that would have to be aligned.

Applications

The most obvious application for collective detection is enabling a navigation fix in circumstances where degraded signals cause traditional acquisition to fail. The sweet spot of collective detection is providing a rapid but coarse position solution in a weak signal environment. The solution can be found in less time because information is evaluated cohesively across satellites. This is especially clear when the algorithm is compared to computationally intensive long integration techniques.

There are several ways that collective detection can support urban navigation. This capability benefits long endurance users who desire a moderate accuracy periodic fix for monitoring purposes. In some circumstances, the user may wish to initiate traditional tracking loops for a refined position estimate. However, if the signal environment is unfavorable at the time, this operation will waste valuable power. The collective detection response indicates the nature of the current signal environment, such as indoors or outdoors, and can inform the decision of whether to spend the power to transition to full GPS capabilities.

In urban applications, deeply integrated GPS/INS solutions tolerate GPS outages by design. However, if the outage duration is too long, the estimate uncertainty will eventually become too large to allow conclusive signal detection to be restored. Running collective detection as a background process could keep deeply integrated filters centered even in long periods of signal degradation. Because collective detection approaches the acquisition problem from a position space instead of the individual satellite line-of-sight space, it provides inherent integrity protection. In the traditional approach, acquiring a multipath signal will pollute the overall position fix. In collective detection, such signals are naturally exposed as inconsistent with the position estimate.

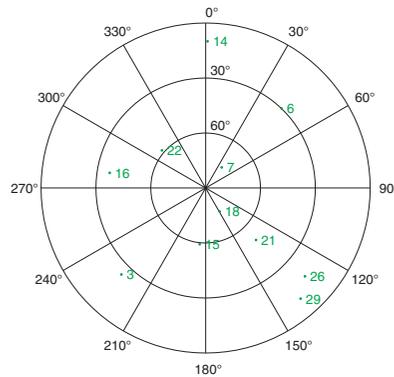
Another use would be to initialize landmark correlation algo-

rithms in vision navigation. Landmark correlation associates street-level video with 3D urban models as an alternative to (GPS) absolute position and orientation updates. This technique associates landmarks observed from ground-level imagery with a database of landmarks extracted from overhead-derived 3D urban models. Having a coarse position (about 100 meters accuracy) enhances initialization and restart of the landmark correlation process. Draper Laboratory is planning to demonstrate the utility of using collective detection to enable and enhance landmark correlation techniques for urban navigation.

In all of these applications, collective detection is straightforward to implement because it simply uses the output of correlation functions already performed on GPS receivers.

Simulations and Processing

The new algorithm has been tested using



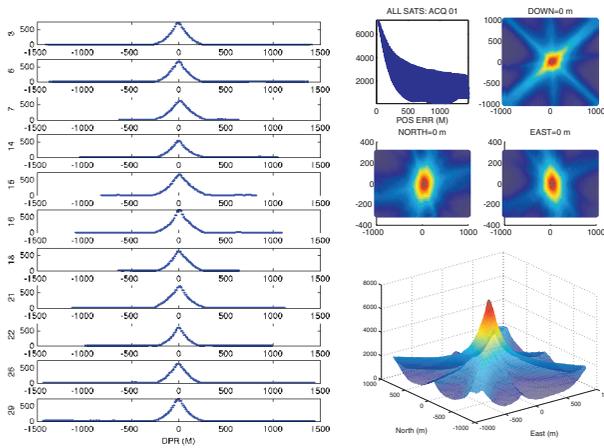
▲ **FIGURE 4** Satellite visibility for simulated data.

live-sky and simulated data collected by a Draper Laboratory wideband data recorder. A hardware GPS signal simulator was used to simulate a stationary observer receiving 11 equally powered GPS signals that were broadcast from the satellite geometry shown in **FIGURE 4**. The data recorder and the signal simulator were set up in a locked-clock configuration with all of the

simulator's modeled errors set to zero. No frequency offsets should exist between the satellites and the receiver. A clock bias, however, does exist because of cable and other fixed delays between the two units. The data recorder houses a four-channel, 14-bit A/D module. It can support sample rates up to 100 MHz. For this work, it was configured to downconvert the signal to an IF of 420 kHz and to produce in-phase and quadrature samples at 10 MHz.

Results and Discussion

To combine satellites, a position domain search space is established, centered on the correct location and receiver clock bias. A grid spacing of 30 meters over a range of ± 900 meters in north and east directions, and ± 300 meters in the vertical. In the first simulated example, the correlation power for all the satellites is summed on the position grid using a single 1-millisecond integration period. In this case, the true



▲ **FIGURE 5** Results from a simulated 11-satellite, 40 dB-Hz scenario. Left panel shows the individual satellite correlations as a function of range error. The upper left plot in the upper right panel shows the combined correlation value as a function of the magnitude of the position error, and going clockwise from there, the projection of the correlogram onto the east-north plane, the north-down, and the east-down planes. The lower right shows a 3D view of the correlogram in the east-north plane.

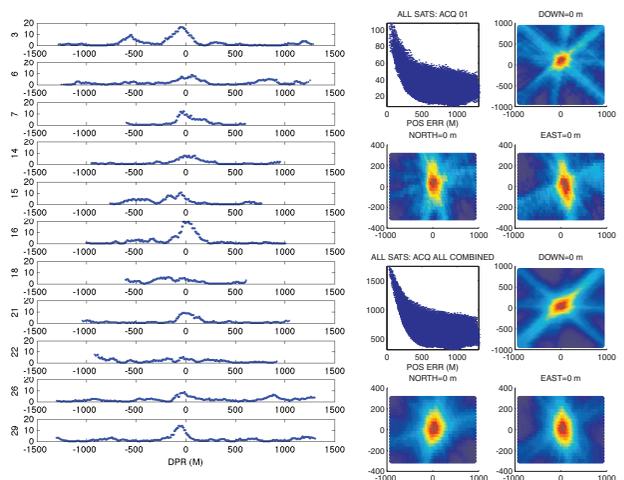
carrier-to-noise-density ratio for each signal is 40 dB-Hz. The results are shown in **FIGURE 5**. The plots in the left panel show the individual signal correlations as a function of range error. The four plots in the upper-right panel show several views of the combined correlation as a function of position error. The upper-left plot in the panel shows the correlation value as a function of the magnitude of the position error. The upper-right plot shows the correlation as a function of the north-east error, the lower-left the north-down error, and the lower-right the east-down error. Notice how the shape of the constant power contours resembles the shape of the constant probability contours that would result from a least-squares solution's covariance matrix. The final plot, the bottom-right panel, shows a 3D image of the correlation power as a function of the north-east error. It is clear in these images that in the 40 dB-Hz case each satellite individually reaches the highest correlation power in the correct bin and that the combined result also peaks in the correct bin. In the combined satellite results, each individual satellite's correlation power enters the correlogram as the ridge that runs in a direction perpendicular to the receiver-satellite line-of-sight vector and represents a line of constant pseudorange.

FIGURE 6 shows a similar set of graphs for a simulator run at 20 dB-Hz. The plots in the left panel and the four plots in the upper-right panel show the individual and combined correlations, as in Figure 5. In the lower-right panel, the 3D image has been replaced with correlations calculated using 20 noncoherent 1-millisecond accumulations. The indistinct peaks in many of the individual correlations (left panel) suggest that these signals may not be acquired and tracked using traditional methods. Those signals, therefore, would not contribute to the navigation solution. Yet in the combined case, those indistinct peaks tend to add up and contribute to the navigation solution. These results indicate

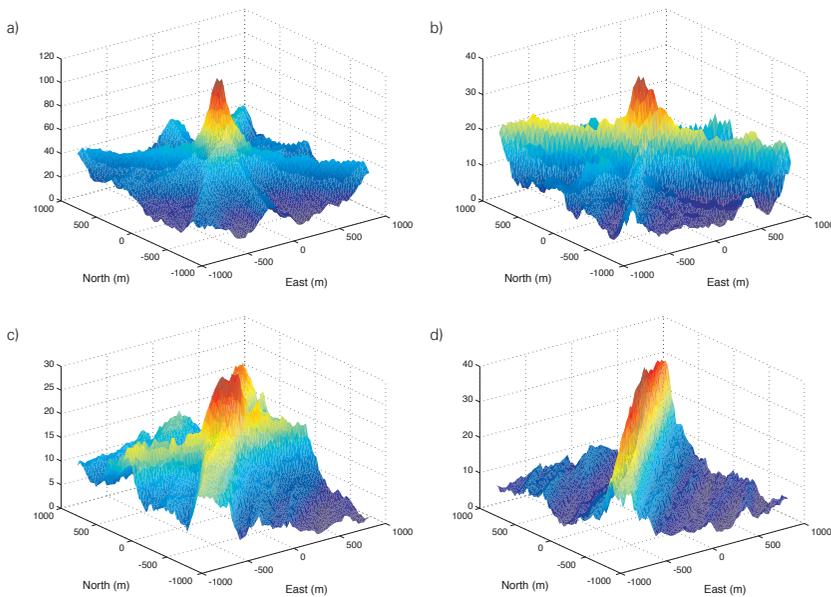
the feasibility of using the information in weak signals that may not be detectable using traditional methods and short acquisition times. The situation is further improved by increasing the number of noncoherent integration periods.

Impact of Reduced Geometry. Of course, it is a bit unrealistic to have 11 satellites available, particularly in restricted environments, so we also considered three subsets of four-satellite acquisitions, under the same signal levels. **FIGURE 7** compares the position domain correlograms for the following 20 dB-Hz cases: (1) a good geometry case (PRNs 3, 14, 18, 26), (2) an urban canyon case where only the highest 4 satellites are visible (PRNs 15, 18, 21, 22), and (3) a weak geometry case where just a narrow wedge of visibility is available (PRNs 18, 21, 26, 29). As expected, the correlation power peak becomes less distinct as the satellite geometry deteriorates. The pattern of degradation, morphing from a distinct peak to a ridge, reveals that the position solution remains well constrained in some directions, but becomes poorly constrained in others. Again, this result is expected and is consistent with the behavior of conventional positioning techniques under similar conditions.

Focusing on Clock Errors. In some real-world situations, for example, a situation where a receiver is operating in an urban environment, it is possible for the position to be fairly well known, but the clock offset and frequency to have substantial uncertainty. **FIGURE 8** shows how the combined satellites approach can be used to improve sensitivity when viewed from the clock bias and frequency domain. The figure presents example 1-millisecond correlograms of clock bias and clock drift for three 20 dB-Hz cases: (1) a single GPS satellite case; (2) a four-satellite, good geometry case; and (3) an 11-satellite, good geometry case. The assumed position solution has been offset by a random amount (generated with a 1-sigma of 100 meters



▲ **FIGURE 6** Results from a simulated 11-satellite, 20 dB-Hz scenario. Same description as for Figure 5, except that the lower-right panel now contains four plots that show the position domain correlogram for 20 noncoherent 1-millisecond accumulations.



▲ **FIGURE 7** The effects of satellite geometry in 1-millisecond integration, 20 dB-Hz cases: (a) east-north position domain correlograms for 11 satellites, (b) for four satellites with good geometry, (c) for four highest elevation angle satellites, and (d) for four satellites with weak geometry.

in the north and east components, and 20 meters in the up component), but no individual satellite errors are introduced. These plots clearly show the improved capability for acquisition of the clock errors through the combining process.

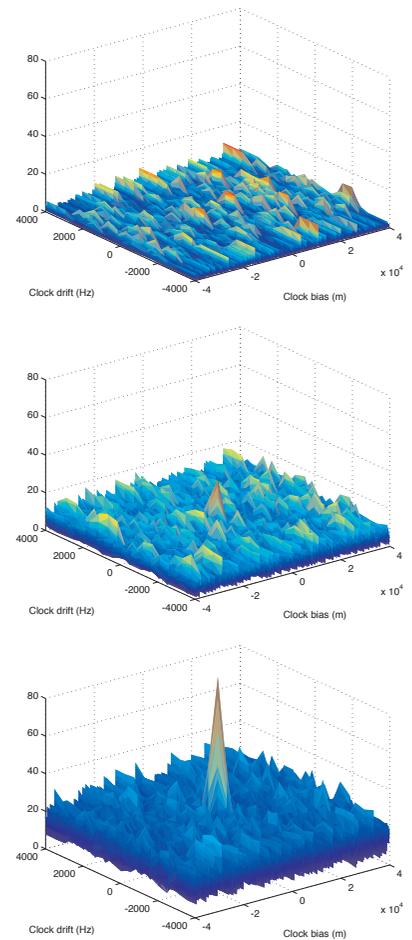
Live Satellite Signals. **FIGURE 9** shows combined correlograms derived from real data recorded using an outdoor antenna. The first example includes high-signal-level satellites with 1.5-second noncoherent integration. The second example includes extremely attenuated satellite signals with a long noncoherent integration period of six seconds.

The plots in the upper-left and upper-right panels show combined correlograms as a function of the north-east position error for satellite signals with carrier-to-noise-density ratios of 48 dB-Hz or higher. The plots in the lower-left and lower-right panels show combined correlograms resulting from much weaker satellites with carrier-to-noise-density ratios of roughly 15 to 19 dB-Hz, using a coherent integration interval of 20 milliseconds and a noncoherent interval of six seconds. **FIGURE 10** shows one of the individual single-satellite correlograms. In this attenuated case, the individual satellite power levels are just barely high enough to make them individually

detectable. This is the situation in which collective detection is most valuable.

Conclusions

The example results from a hardware signal simulator and live satellites show how the noncoherent combination of multiple satellite signals improves the GPS position error in cases where some of the signals are too weak to be acquired and tracked by traditional methods. This capability is particularly useful to a user who benefits from a rapid, but coarse, position solution in a weak signal environment. It may be used to monitor the quality of the signal environment, to aid deeply coupled navigation, and to initiate landmark recognition techniques in urban canyons. The approach does require that the user have some *a priori* information, such as a reasonable estimate of the receiver's location and fairly accurate knowledge of the GPS ephemerides. Degradation in performance should be expected if the errors in these models are large enough to produce pseudorange prediction errors that are a significant fraction of a C/A-code chip. Absent that issue, the combined acquisition does not add significant complexity compared to the traditional approach to data processing. It can be used to enhance performance

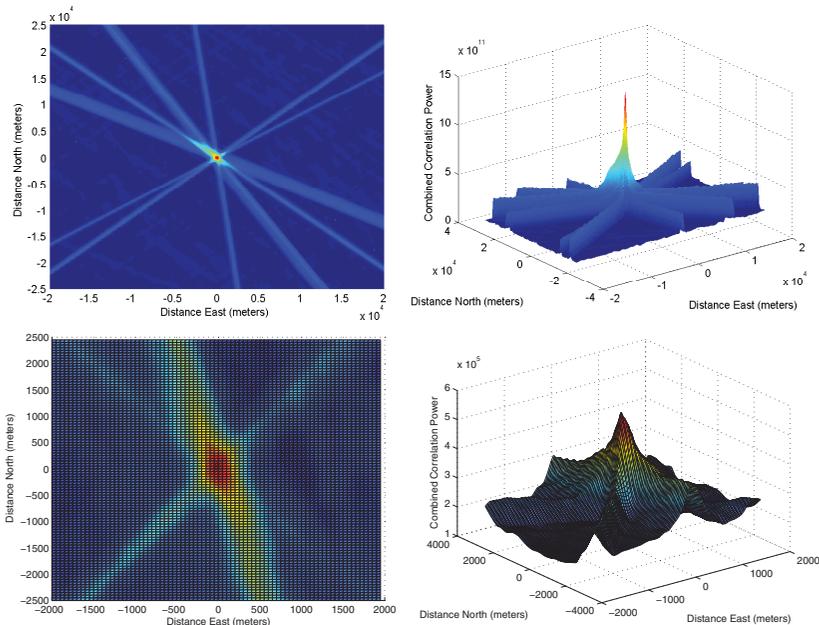


▲ **FIGURE 8** Correlograms in clock bias and drift space, with random position error (1-sigma of 100 meters in north and east, 20 meters in vertical). The top plot is for a one-satellite case; the middle is for a four-satellite, good geometry case; and the bottom is for an 11-satellite, good geometry case. All calculated with a single 1-millisecond integration period.

of existing acquisition techniques either by improving sensitivity for the current noncoherent integration times or by reducing the required integration time for a given sensitivity. Further development and testing is planned using multiple signals and frequencies.

Acknowledgments

The authors appreciate the contributions of David German and Avram Tewtewsky at Draper Laboratory in collecting and validating the simulator data; Samantha Krenning at the University of Colorado for assistance with the simulator data analysis and plotting; and Dennis Akos



▲ **FIGURE 9** Horizontal position domain correlograms for rooftop data using 20-millisecond noncoherent integrations. The top two plots are for a strong signal case (48-58 dB-Hz) with 1.5-second noncoherent accumulations. The bottom two plots are for a weak signal case (15 to 19 dB-Hz) with 6-second noncoherent accumulations.

at the University of Colorado for many helpful conversations and for providing the Matlab software-defined radio code that was used for setting up the acquisition routines. This article is based on the paper “Enhancing GNSS Acquisition by Combining Signals from Multiple Channels and Satellites” presented at ION GNSS 2009, the 22nd International Technical Meeting of the Satellite Division of The Institute of Navigation, held in Savannah, Georgia, September 22–25, 2009.

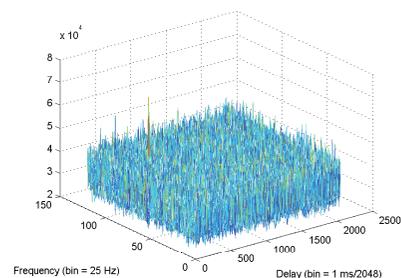
The work reported in the article was

funded by the Charles Stark Draper Laboratory Internal Research and Development program.

Manufacturers

Data for the analyses was obtained using a **Spirent Federal Systems** (www.spirentfederal.com) GSS7700 GPS signal simulator and a **GE Fanuc Intelligent Platforms** (www.gefanuc.com) ICS-554 A/D module. 🌐

PENNA AXELRAD is a professor of aerospace engineering sciences at the University of



▲ **FIGURE 10** Single satellite, weak signal (15-19 dB-Hz) correlogram for rooftop data, 20-millisecond coherent integrations, 6-second noncoherent accumulations.

Colorado at Boulder. She has been involved in GPS-related research since 1986 and is a fellow of The Institute of Navigation and the American Institute of Aeronautics and Astronautics.

JAMES DONNA is a distinguished member of the technical staff at the Charles Stark Draper Laboratory in Cambridge, Massachusetts, where he has worked since 1980. His interests include GNSS navigation in weak signal environments and integrated inertial-GNSS navigation.

MEGAN MITCHELL is a senior member of the technical staff at the Charles Stark Draper Laboratory. She is involved with receiver customization for reentry applications and GPS threat detection.

SHAN MOHIUDDIN is a senior member of the technical staff at the Charles Stark Draper Laboratory. His interests include GNSS technology, estimation theory, and navigation algorithms.

MORE ONLINE 
Further Reading
 For references related to this article, go to gpsworld.com and click on Innovation under Resources in the left-hand navigation bar.

AD & COMPANY INDEX

ADVERTISERS

COMPANY	PAGE
Anatech Microwave.....	32
IFEN GmbH.....	2
JAVAD GNSS.....	11–18
Navcom.....	67

COMPANIES

Avera.....	34
DigitalGlobe.....	24
GE Fanuc Intelligent Platforms.....	64
GPS Tuner.....	24
Hemisphere GPS.....	27
Hirschmann Car Communication.....	65
Jackson Labs Technologies.....	65
John Deere.....	65

COMPANY	PAGE
NovAtel.....	9, 68
OmnISTAR.....	25
Racelogic.....	31
Septentrio.....	7

Microsoft.....	24
Navteq.....	24
Networks in Motion.....	25
Racelogic.....	34
Sierra Nevada Corporation.....	25
Spirent Federal Systems.....	64
Surrey Satellite Technology Ltd.....	25
Tele Atlas.....	24

COMPANY	PAGE
746th Test Squadron.....	33
Spirent Federal.....	5
Topcon.....	27
Trimble.....	23

Telecommunication Systems.....	25
Third Eye Maps.....	27
3-GIS.....	65
Topcon.....	66
Trimble.....	26

More companies appear in the Receiver Survey, pages 35–57.