

Don't Dump That Data!

Thorough Use of GNSS Measurements



JAMES L. FARRELL
VIGIL, INC.

A key step toward maximizing robustness in GNSS signal processing is to abandon habitual practices that offer convenience or familiarity, but are not essential, including the discarding of some types of data. Myriad cases have been documented of signal degradation from such phenomena as interference, masking, attenuation, ionosphere scintillation, and multipath. Using flight-validated experience, this article describes methods for extracting maximum achievable performance in the presence of unrelenting and severe adverse effects in the operational environment.

© iStockphoto.com/Christoph Emel

A constantly recurring refrain within the satellite navigation community is the self-evident observation that, with more satellites, constellations, and frequencies appearing over the next few years, GNSS performance will improve.

Until those expectations become reality, though, another major opportunity is being almost universally overlooked: making *far* more efficient use of all the GNSS measurement data we have available right now.

Even after addition of more satellites and frequencies, a need for full extraction of information will continue in many operations. Even with more satellites and signals, users will still encounter various sources of signal degradation (e.g., intentional and/or unintentional interference, multipath, ionospheric

disturbances), signal blockage (obstruction by foliage, terrain masking, or even elements of the structure carrying the antenna), and attenuation of already weak GNSS signals in sensitive (e.g., indoor) operations.

To those eagerly awaiting the new signal resources, claims of major improvement now — from existing data without further breakthroughs — may provoke immediate skepticism. Those willing to pursue the issue, however, can realize those benefits even before the new signal resources come on line.

In this short article, I describe how satellite signal data unexploited by extant user equipment, can actually be put to work. I offer the usage of sequential changes in carrier phase data, including those from low-elevation GNSS satellites whose signals are omitted from receiver

processing because of concerns over atmospheric propagation effects.

Velocity from Carrier Phase Changes — Be Thorough!

Before plunging into a full discussion of the topic, here's a preview. Sequential changes in GNSS carrier phase, over one-second intervals, are

- typically accurate to within a centimeter
- immune to cycle count uncertainties — the integers cancel in the subtraction forming the one-second difference
- amenable to treatment as if they were independent — and therefore — (next item)
- *not* required to be continuous (We routinely cope with gaps in a data stream.)

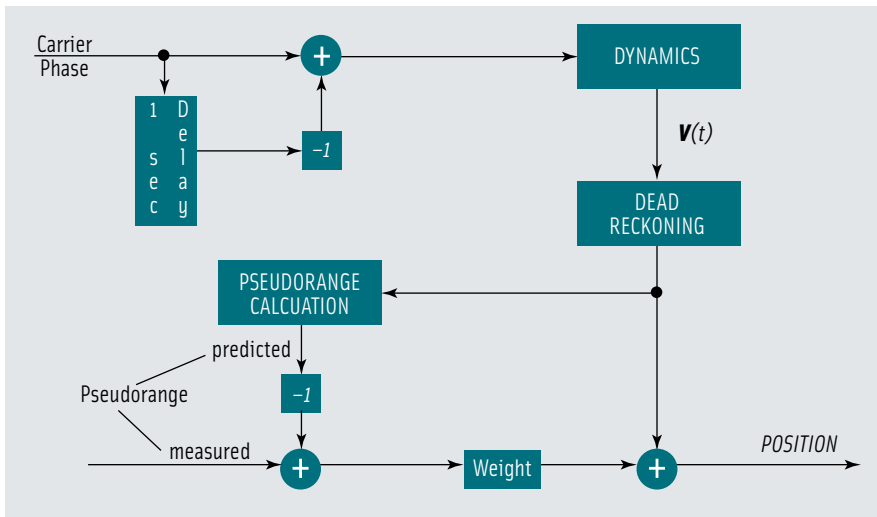


FIGURE 1 Conceptual design for tracking sequential changes in carrier phase

- so insensitive to effects of the ionosphere/troposphere that *no mask angle is needed for those* — (again, benefiting from cancellation of the unknowns) — and
- therefore, useful as a precise and dependable source of velocity history.

Velocity histories thus derived feed directly into dead-reckoning, and pseudoranges (protected by mask as usual and edited via receiver autonomous integrity monitoring (RAIM) and other signal processing techniques) prevent unattended growth in the position error. An inertial measurement unit (IMU) is not required; this can be accomplished with or without one.

Assertions made in this article are based on flight-validated experience. Rather than repeatedly referring to test results, they are mentioned only in passing where needed for clarification.

All derivations, details, steps, answers to accompanying questions/correlation effects/issues/paradoxes and so forth — *as well as* supporting test results — are documented in my recently published book, *GNSS Aided Navigation and Tracking: Inertially Augmented or Autonomous*, cited in the Additional Resources section.

This article is intended to provide only a simple description of this approach with no mathematical development.

Figure 1 shows a simple system driven by only GNSS data (no IMU nor any

other sensor). Outputs are geographic velocity and position as usual. Velocity is determined from sequential changes in carrier phase over one-second intervals. A continuous time history of velocity is simply integrated to produce position.

To counteract the growth of position error due to this integration, only pseudoranges, (i.e., no carrier phase information) are used for position determination, and, conversely, no pseudorange information is used for calculating velocity. The dynamic model merely consists of a constant acceleration vector.

To blunt the effect of modeling discrepancies from acceleration changes, the data-averaging durations are kept short (i.e., a few seconds). In analogy with spline fitting, then, estimates continuously evolve — formed as combinations of current and recent observations. That latter concept has solid support from a long history of success in myriad tracking applications.

Unlike virtually all existing operational methods, then, this approach makes full use of the carrier phase, even if permanently ambiguous *and* intermittent. That latter trait — intermittency — is made acceptable by reason of another unusual feature reflected in the design of a GNSS receiver developed at The Ohio University Avionics Center and described in the article by F. van Graas et alia cited in Additional Resources.

This OU receiver design replaces correlators and tracking loops with fast

Fourier transform (FFT) processing, which offers

- unconditional stability (in contrast to a conditionally stable third-order loop)
- unconditional access to all inverse FFT cells, allowing a broader effective correlation search than the narrow subset that a tracker would cover
- linear phase shift (as opposed to non-linear phase of a track loop, causing group delay variations).

In addition, the user can easily be empowered to select the durations of both the GNSS carrier difference and code data blocks. That option provides full flexibility and capability for versatile operation.

Obstacles to Adoption

The approach just presented is so effective and has such simple dynamics that it prompts a question: Why has this capability remained unused for all these years? The answer to this question has multiple parts.

Selective Availability. A derivation relating sequential changes in carrier phase to precise velocity, accounting for the large excursions of satellites over the same interval, never appeared until shown in my previously cited book. For many years, the presence of the selective availability (SA) time dither limited the motivation to investigate this approach. Despite an obvious advantage (i.e., cancellation of unknown cycle counts), which was noted in my 2001 ION/CIGTF presentation (see Additional Resources) — again using data collected from before SA's removal — stand-alone capability wasn't possible.

Land trials in a van—conducted under SA conditions but employing these principles — were successful, but they depended on receiving corrections from a ground station. However, differential GPS performance was already brilliant without the programming-intensive processing required for converting sequential carrier phase changes into velocities; so, the scheme did not find wide acceptance.

Legacy of Past. Technical obstacles that shaped receiver designers' thinking in the past may blind them to the practicality of such methods in today's world. The approach discussed in this article requires added processing easily possible with today's electronics.

Position Fixation. The industry has long prioritized instantaneous position accuracy above other considerations. As a figure of merit, instantaneous position-fixing seems to offer — at least conceptually — an easily understood basis for decisions. This is valued especially in applications involving certification, wherein anything not immediately apparent can be open to question.

For certification that stance may or may not ever change but, in any case, the following brief discussion of alternative decision criteria is in order.

Among the most obvious reasons for broader figure-of-merit criteria: many operations are not rigidly governed by an inescapable need to substantiate the minimization of theoretical risk. Furthermore, whether operations are required to quantify risk or not, "accuracy" is widely understood to have multiple meanings because the concept of "error" has multiple interpretations (e.g., RMS, containment, CEP, SEP, and so on).

A corresponding variety of error metrics is needed to accommodate different conditions and/or performance levels — as well as differences in context (absolute, repeatable, relative). The latter item in particular can, in some applications, impose additional demands. For instance, a need may arise for adherence to a uniform datum reference because, if coordinates are being subtracted to yield relative locations, misunderstanding arising from nonuniformity will undercut the validity of the results.

Finally, the value of reducing instantaneous position error to ever-smaller amounts — centimeter? millimeter? less than one millimeter? — should be judged in light of diminishing returns. For example, even if we could pinpoint the instantaneous location of the antenna phase center on a supersonic aircraft to within a micron, our efforts would

be better spent on determining future position.

Freedom from commonly imposed constraints, ironically, can often produce better prospects of success — just by using information that tightly controlled systems would discard. And that immediately raises another issue, the safety of using such information, which we will address next.


Making Previously Rejected Data Safe

In addition to obvious features in Figure 1 (segmentation and sequential change in carrier phase over the course of one second), another procedural step is fundamental to the approach: single-measurement RAIM, with optional subsequent verification by more familiar multi-satellite RAIM methods. Thresholds for each residual can be set to control probabilities of alarm and missed detection, with the same rigor used for deriving RAIM tests — but with one significant improvement. Each separate observation can be weighted according to its individual credibility (e.g., based on signal strength, elevation angle, and so on).

Furthermore, rather than selecting only the space vehicles (SVs) offering the best RAIM geometry, every satellite in view can be acceptance-tested as a candidate for insertion into estimation processing. Optionally a designer might choose to employ additional well-known procedures such as across-SV differencing (followed by whitening to undo resulting correlations) to remove user clock errors.

The measurement separation — between pseudoranges and carrier phases — just described is total: a pseudorange rejected from a specific satellite does *not* affect decisions about whether to accept that same satellite's carrier phase sequential-change data.

Of course, mask angle is prescribed to reject pseudoranges from low-elevation satellites in order to edit data containing long ionospheric and tropospheric delays. When tracking sequential change in carrier phase, however, these same SVs should not be automatically rejected.



Over a one-second interval, those sequential phase changes are generally accurate to within a centimeter RMS. If multipath or ionospheric disturbances degrade some of them, data editing (via both single-measurement and subsequent multi-satellite RAIM) can exclude them. *Some* of the data at an instant can be rejected without rejecting *all* of the data at that instant. Especially when a shortage of data exists, use of partial information may be vital for more than one reason. Not only does it add measurements that would otherwise be ignored, but it can improve coverage geometry, which may be pivotal for a successful application.

Successful flight test results described in *GNSS Aided Navigation and Tracking: Inertially Augmented or Autonomous* were obtained after SA removal; that was stand-alone operation. What was previously a limited range of application — operation with a ground station, already brilliantly successful without the capabilities described herein — now expands to unrestricted usage.

An opportunity for major performance improvement is therefore available for general operation. Hardly any in the industry are taking advantage of this at present but that is expected to change; the cost is too low while the benefit is too high to remain untapped.

Conclusion

My discussion here was intentionally limited to fundamental considerations and included only a restricted subset of applications. A wider scope of issues (e.g., inertial updating, free-inertial coast, tracking, etc.) can be found addressed at length in my book.

With the welcome addition of more satellites and more frequencies, these considerations will continue to be applicable, for multiple reasons. Here are some of those reasons:

- Hindrances to signal detection, previously discussed, will remain.
- Interoperability is most challenging for items with highest accuracy expectations. With carrier phases from different constellations, sequential changes are far easier to

mix than the phases themselves.

- A microelectromechanical system (MEMS) IMU has much less free-inertial coast capability than a navigation-quality INS, but demands to support GNSS-era performance (centimeters per second rather than yesteryear's nautical miles per hour) will endure.
- Continued growth, often gradual but inexorable, will occur in demands for additional capabilities.
- Those demands will be met by increased sophistication in knowledge from continued research.

Additional Resources

[1] Farrell, J. L., "Carrier Phase Processing Without Integers," pp. 423–428, *Proceedings of Institute of Navigation 57th Annual Meeting/CIGTF 20th Biennial Guidance Test Symposium*, Albuquerque, New Mexico, June 11–13, 2001

[2] Farrell, J. L., *GNSS Aided Navigation and Tracking: Inertially Augmented or Autonomous*, American Literary Press, NavtechGPS (distr.), 2007

[3] van Graas, F., Soloviev, A., Uijt de Haag, M., Gunawardena, S., and Braasch, M., "Comparison of two approaches for GNSS receiver algorithms: batch processing and sequential processing considerations," ION GNSS-2005.

Author



James L. Farrell is the principal in Vigil, Inc., a senior member of IEEE, and a registered professional engineering consultant in Maryland. He received a Ph.D. from the University of Maryland.

His professional career includes two years each at Honeywell and Bendix-Pacific, plus 31 years at Westinghouse in design/simulation/validation of estimation algorithms for navigation and tracking, fire control system design, test data generation for bench validation, and INS update and transfer alignment algorithm design. During that time, Farrell worked on program development for USAF-WPAFB (F/C evaluation) and NASA (estimation of satellite attitude, orbit determination) as well as missile guidance optimization and digital communication design. He is author of *GNSS Aided Navigation and Tracking* (2007; NavtechGPS distributor), *Integrated Aircraft Navigation* (Academic Press, 1976). Active in RTCA for several years, he served as co-chairman of the GPS Integrity Fault Detection and Isolation Working Group within SC-159. He has taught various tutorials and courses at Marquette and UCLA. 

see what's new
online

insidegnss.com