Robust Design for GNSS Integration

James L. Farrell, *VIGIL Inc.*

Presented at Institute of Navigation GNSS 2008

Savannah GA

BIOGRAPHICAL

James L. Farrell (Ph.D., U. of MD `67; former ION Air Nav representative, senior member of IEEE, former AIAA local board member, member of τβπ, ηκν, πμϵ) is a registered professional engineering consultant in Maryland. Work includes teaching at Marquette and UCLA, two years each at Honeywell and Bendix-Pacific, plus 31 years at Westinghouse in design/simulation/ validation of estimation algorithms for navigation and tracking [*e.g.*, F16 AFTI, B1 radar, SDI; fire control system design, test data generation for bench validation; INS update and transfer alignment algorithm design; program development for USAF-WPAFB (F/C evaluation) and NASA (estimation of satellite attitude, orbit, deformation); missile guidance optimization, MLE boundaries] plus digital communication design (sync, carrier tracking, decode). He is author of *GNSS Aided Navigation and Tracking* (2007; NavtechGPS dist.), *Integrated Aircraft Navigation* (Academic Press, 1976; now in paperback after five printings), and of various columns plus over 80 journal and conference manuscripts. Active in RTCA for several years, he served as co-chairman of the GPS Integrity Fault Detection and Isolation Working Group within SC-159.

ABSTRACT

Loss of satellite navigation data can be catastrophic in many operations. Common mechanization approaches often fail to provide adequate protection under adverse conditions. Remedial methods are largely based on refocusing of objectives – *e.g.*,

- aircraft with dozens of meters wingspan don't have to be located to within *1*-cm as they move at 400 kts,
- In most operations, *dependably* providing continuous position to within a few meters is highly preferable to *usually* positioning within ± *1* cm.

Means for full usage of partial data are more extensive in scope than generally realized – and some that are known have only recently become feasible through technological advances. Flight-validated methods described herein, inside and external to the receiver, remain applicable with or without IMU augmentation.

The author is indebted to Ohio Univ. and FAA for support of this work.

INTRODUCTION

Conventional GPS processing and integration methods have performed successfully for reasonably benign conditions. There are, however, myriad documented cases of degradation from various effects including intentional and/or unintentional interference, masking, attenuation, ionosphere scintillation, obscuration, multipath, etc. With multiple independent phenomena occurring in tandem, the result can be repeated breaks in phase tracking and inability to resolve, with sufficient confidence, cycle ambiguities. The degradations affect pseudoranges (thus code tracking) as well. Even when some track loops function with some continuity, geometry of those SVs being tracked can fail to support RAIM. The goal of this paper is to describe numerous flight-validated methods devised to extract maximum achievable performance in the presence of unrelenting severe adversity.

Many situations that defeat current mechanizations are by no means hopeless. One key to success under adversity is utilization of partial – even fragmented – information. Conventional approaches depend heavily on some subset or some combination of the following:

- high-order track loops fed by correlators
- resolution of carrier phase ambiguities
- full-fix with acceptable GDOP
- fixes with RAIM-enabling geometry

A key step toward maximizing robustness is to abandon habitual practices that offer convenience of familiarity, but are not essential. None of the four features just noted are essential for configurations described herein. While limiting situations can always be conjured up to defeat any man-made design, less extreme but still taxing conditions can be surmounted by going beyond conventional methods. Many existing procedures impose needs inherited from or traceable to early mechanizations. Digitization long ago enabled removing many of the resulting limitations.

The best humanly possible strategy is to deliver whatever performance is reachable from all available information, incomplete as it may be. For decades it has been feasible to combine intermittent partial data – of different types at varying accuracies with different sensitivities from different directions at different times – and extract all benefit offered. Decisions adopted herein offer no more and no less than that.

As is often the case, a full bibliographical citing would be a daunting task. Instead, minimal references are given here, with optional further investigation available from references cited therein. It can be seen from those inner citings that

- methods described in [1] and [2] were devised and documented years earlier than those publications, and
- all methods were carried beyond development, to verification in testing. Not all techniques were present together in the same tests (*e.g.*, GPS/INS flight tests of [2] used receivers with interfaces allowing tight – but not ultratight – coupling). Still, state-of-the-art performance was realized with low-cost equipment – and, with future interfaces giving unrestricted access, methodology discussed here can be exploited still further.

For techniques to be recommended, descriptions offered here emphasize the advantages while avoiding repetition of details documented in [1] and [2]. The primary aim is not presentation of theory but advocacy for greater attention to robustness in configuring future systems for use under adversity.

APPROACH

The next four subsections define procedures for release from operational constraints to allow maximum usage of partial information. Some of the methods just capitalize on recent technological improvements that enable exploitation of older principles. Some of the theoretical concepts were developed in ways differing from convention; the intent was to use any means of enhancing ability to carry on under challenging conditions.

The first two subsections describe test-validated ways to alleviate needs for tasks involved in customary usage of GPS carrier phase. One task (maintenance of continuous tracking with no undetected cycle slips) is an essentially omnipresent goal in operational mechanizations that employ carrier loops. A second task (cycle count ambiguity resolution), though not quite universal, is widespread; ambiguous carrier phase is used in only a minority of operations.

Adherence to demands imposed by these tasks has provided considerable success with benign conditions. Those demands are burdensome, however, with carrier phase under siege by masking, attenuation, intentional and/or unintentional interference, multipath, obscuration, ionospheric scintillation, etc. There have been abundant reported instances wherein the demands could be met only marginally, or with considerable difficulty, or not met satisfactorily. Highest priority is then given to identifying means for continued operation without attempting to satisfy unachievable nonessential requirements.

The latter two subsections expand on usage of partial fix data. Tight coupling in GPS/INS integration, with across-SV differencing for user clock error removal, is only the beginning. Separate processing is employed for dynamics based on carrier phases and position based on pseudoranges, with pre-whitening to account for correlation effects. Follow-through then extends to integrity testing applied separately to each individual measurement, producing results in full correspondence with widely used rigorous matrix decomposition methods. Correlation effects are also addressed in defining the integrity criteria and, wherever geometry permits, multi-SV RAIM retesting is retained as an option.

Replacement of Correlators and Track Loops

For about a century it has been known that convolution in the time domain corresponds to multiplication in the frequency domain (and vice-versa). Application to GPS receiver implementation [1] (a relatively recent happening, with today's processor speeds and A/D converters) removes three fundamental limitations shared by even the most ingenious poles-and-zeros transfer functions conceivable:

- restriction to a narrow subset of cells
- conditional stability
- group delay variation among spectral constituents of the signal

Although stability of a third order track loop is often very good, it can fail under adverse conditions – there can be times when the signal being sought resides beyond the cells being tracked. By contrast, the approach used in [1] enables unconditional access to all FFT cells, with uniform group delay across the signal spectrum (the linear phase shift feature follows from the FFT's all-zero trait). Even without further benefits of the approach described in [1] (unique means of code stripping, etc.), it clearly offers superior capability. Combined with the segmented GNSS/IMU configuration described in the next subsection, it also offers opportunities to supplant ultra-tight coupling and all lesser methods of integration.

Retention of Carrier Phase Ambiguities

Both the benefit (sub-wavelength accuracy) and the pitfall (vulnerability to false success indications) of cycle ambiguity resolution are universally recognized. Complete elimination of the risk results from employing ambiguous carrier phase throughout. While the operation described in [2] is not the only instance of applying ambiguous carrier phase, its subtle differences from other approaches provided significant refinement (*e.g.*, RMS leveling errors at a few tenths mrad, not tenths of a degree). To ensure clarity of the advantages offered, main features are delineated here.

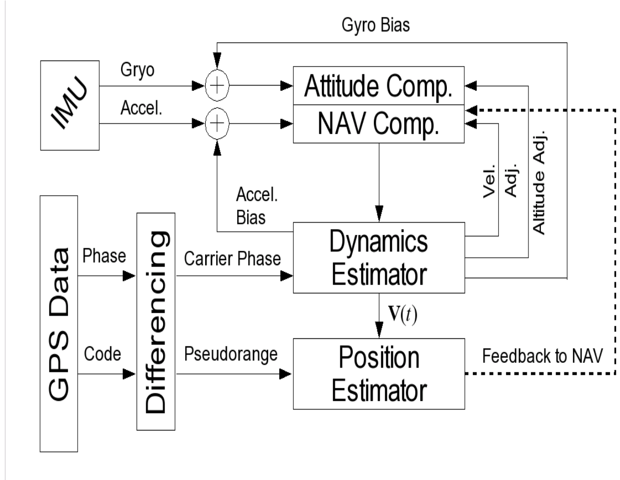

Figure 1: Segmented Block Diagram

Figure 1 matches Figure 5.1 of [2], but with improved placement of tags for velocity and attitude adjustments (*i.e.*, with verbage closer to corresponding lines). The main departure from conventional GPS/INS integration is segmentation. Sequential changes in carrier phase can adjust dynamics only, with velocity history feedforward to the position estimator for integration into *a priori* position, adjusted by pseudoranges only. Removal of risk from cycle count uncertainty is immediately apparent; unknown integers cancel in subtraction to form the sequential phase changes. Slightly less obvious – but potentially more important – is prompt recovery from interruption in carrier tracking. After loss of SV data continuity, adjustments in dynamics can resume as soon as the difference in two successive phase measurements passes a rigorously derived integrity test (also discussed in a later subsection here and fully documented in [2]). With an unknown jump in the unknown cycle count, then, the dynamics estimator experiences only a short gap in its incoming data stream, easily accommodated.

Various additional points, not all of them obvious, are worth noting. For brevity only the main issues are given here; again, formulations and details are documented in [2].

- An alternative to segmentation is of course the common configuration with position and all higher order states together in one estimator. Optionally in that case, position can be absolute while ground reference stations can be used to facilitate ambiguity resolution. Another common practice employs incremental rather than absolute position, allowing easy initialization while also providing immunity to cycle count bias. Only that latter choice (incremental rather than absolute position) provides that bias immunity, ***and the protection stops there***. A change in an unknown cycle count error ***violates that*** alternative model; not only the incremental position but velocity and all higher order states become questionable until stable operation is restored.
- A similar consideration applies when, instead of pseudorange, unambiguous carrier phase updates position, giving precise (sub-wavelength) location. A conventional estimator loses credibility of *all* states (not only position) when cycle counts become discontinuous. The segmented estimator, easily accommodating gaps in its incoming data stream, therefore has an important advantage – suppression of transients in dynamics – over conventional methods even with ambiguity-resolved carrier phase.
- Initial reflection might suggest omitting across-SV differencing for the dynamics segment if the user clock can be guaranteed stable for brief differencing intervals. A safer requirement would be user clock drift < 1 cm / s over the effective "memory" (data-averaging duration) of dynamics estimates. When across-SV differencing is retained throughout (as it was in the flight test data processing reported in [2]), the resulting positive correlations among observation errors are taken into account by a simple pre-whitening algorithm.
- In addition to the effect just noted, observation errors for the dynamics segment have negative correlations due to sequential differencing. Resolving that issue necessitated a far deeper theoretical investigation but, for *1*-sec differencing intervals, mechanization and programming did not have to be broadened.
- Sequential differencing offers ease of interoperability with carrier phase data from different constellations. Satellite mislocation will not be sufficient to seriously degrade phase changes over *1*-second and, clearly, it matters little whether those changes came from satellites belonging to GPS or any other group.
- One antenna suffices for GNSS/INS integration in terrestrial operations. Extension is easily made to include dual or multiple antennas for improved azimuth or for space applications.
- Sequential change in distance between a satellite (at present position vector $\mathbf{A}$ with excursion $\mathbf{a}$ over the past second) and a user (at present position vector $\mathbf{R}$ with excursion $\mathbf{r}$ over the past second) is formed from a largely unused but powerful relation :

$$|\mathbf{A} - \mathbf{R}| - |(\mathbf{A} - \mathbf{a}) - (\mathbf{R} - \mathbf{r})|$$

Taylor series expansion produces easily manageable expressions for residuals while suppressing numerical error, and sensitivities (**H**-matrix elements) are formed by expressing $\mathbf{r}$ in terms of velocity and all error states. The fully documented development allows separate handling of propagation and lever-arm effects, including unusually detailed accounting for the latter. Page 104 of [2] provides a table showing all significant terms for one set of difference measurements during a DC3 flight with severe vibration. Algebraic sums of terms, with magnitudes up to several hundreds, all produce residuals of

- either zero or ± 1 to the nearest cm in the table
- 1 cm RMS for almost an hour of flight.

Full Usage of Partial Fix Information

Release from undue dependence on GDOP of synchronized measurements, by optimally weighting each individual observation, has been recognized for many years. The resulting advantages have long benefited GPS/INS in tight coupling. There is less awareness of similar opportunities in using satellite data without inertial information. Section 8.1.2 of [2] shows flight results from unaided GPS data in a segmented estimator configuration based on short-term quasistatic acceleration. Velocity is naturally less accurate than the GPS/INS results (on the order of dm/sec rather than cm/sec), but that algorithm makes the most of unaided data – *sans*-IMU – in the high-vibration environment.

Separate usage of each SV observation, for carrier phases and for pseudoranges independently, was already clear from Figure 1. That separation and independence were likewise employed in the unaided configuration just described. Again, across-SV differencing was used in processing of all flight data but

- all correlation effects were addressed and resolved, by methods fully documented.
- omission of across-SV differencing is potentially permissible for *1*-s carrier phase sequential changes, as discussed in the preceding subsection.

Integrity Test Separation

Usage of individual measurements without first testing for snapshot RAIM might seem to incur undue risk. Complete protection is provided, however, in multiple ways:

- the ratio of each residual to its RMS value is easily computed and tested against a threshold. Thresholds and detectable biases are computed by the same criteria used in the widely accepted parity method. Sections 6.3 and 6.D of [2] demonstrate rigorous conformance to both Kalman estimation and parity equations.
- in any instance where geometry supports RAIM snapshot testing, multi-SV data can be re-entered into parity test; nothing is lost.
- when RAIM is supported by the geometry and snapshot validation is desired prior to acceptance of any individual measurement, the separate observation data can be held, before entry into the estimator, in a one-second buffer until completion of multi-SV RAIM verification.

Processing approaches given for both scalar (individual) and vector (concurrent) data offer unusual advantages not available with methods customarily used –

- Separate processing of each observation of course enables usage of all available data, not just subsets chosen for geometry. It also allows assigning different RMS values to data affected by different conditions (signal strength, elevation angle, etc.).
- For geometry supporting RAIM validation, algorithms given in [2] offer a subtle advantage for navigation. Across-SV differencing allocates all available information for position, velocity, ⋯ – and none for user clock (timing, not in itself important for navigation, is cancelled out of the formulation). With overdetermination, expending nothing on the user clock offers a slight statistical advantage to navigation – the pool of information is shared by fewer states.

For the latter case, again all correlation effects were addressed and resolved. That multi-SV processing encompasses strange matrix forms, but it is important to note: all the complexity was in the derivation. Operational usage requires only a normalizing premultiplication by a known constant matrix. Once that simple computation is performed as shown in [2], the procedure is no more demanding than carrying out the usual parity function.

SUMMARY

An extensive set of flight-validated methods has been discussed, offering ways to survive adverse conditions with no significant cost burden. It is readily acknowledged that these methods should not automatically be followed by all. Loss of data in some operations may be a minor inconvenience. Proof-of-concept experiments, for example, can be scheduled for benign conditions and, even if unexpected challenges arise, procedures can be repeated. Those situations need little or no effort to provide robustness. In many other operations, however, extended periods without satellite data could be catastrophic. Ironically the absence of data often results from unnecessary conditions imposed by popular habits. Many of those habits, traceable to yesteryear's technology limitations, are so ingrained that little is known of means going beyond their reach. This paper offers raised awareness of stronger capabilities and identifies the means.

REFERENCES

1. vanGraas, 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.
2. Farrell, James L., *GNSS Aided Navigtion and Tracking*, American Literary Press, 2007 (NavtechGPS dist.).