Integrity for Tightly Coupled PPP and IMU

Kazuma Gunning, Juan Blanch, Todd Walter; *Stanford University*
Lance de Groot, Laura Norman, *Hexagon Positioning Intelligence*

ABSTRACT

Precise Point Positioning (PPP) offers high accuracy, global positioning, and there is growing enthusiasm for the application of PPP techniques to safety critical systems [1], [2]. We have shown that PPP, in conjunction with techniques developed for integrity in aviation, can be used to produce meter-level protection levels for static, automotive, and flight scenarios [3]. Solution separation, which requires the creation of solutions that are tolerant to faults or sets of faults, has been shown to provide integrity to the PPP algorithm [3]. This paper builds off of previous work to improve a PPP with solution separation algorithm by using an inertial measurement unit (IMU) to improve solution and protection level continuity. Grouping of fault modes, in conjunction with the IMU and PPP, is shown to reduce computational load while maintaining high levels of performance. Finally, this paper introduces a method to reduce the fault exposure time for sequential filters with solution separation.

INTRODUCTION

PPP alone suffers from the limitations of GNSS in that obstructions, which are common in suburban and urban environments, lead to loss of measurements and thus a loss of solution. For PPP, this penalty is particularly severe, as the convergence time can be on the order of tens of minutes. Inertial measurement units (IMU) are incorporated into the PPP with solution separation algorithms in order to increase not only the accuracy but the availability and continuity of the protection levels. Accelerometers and gyroscopes allow the position solution to be propagated even when GNSS measurements are not available, whether that is between GNSS measurements (which are typically output at lower rates than IMU measurements) or during a GNSS outage due to obstruction. When there are short outages, the continuous IMU measurements allow for rapid reconvergence of the PPP solution and, as a result, the protection levels offered by the PPP with solution separation algorithm.

The second contribution in this paper is the grouping of multiple fault modes in conjunction with the IMU and PPP estimator. Grouping of fault modes has been explored previously [4], but with GNSS-only, the weakened geometries in more hostile signal environments can lead to significant performance degradation with respect to the protection levels available. This paper demonstrates that in the inclusion of the IMU alongside fault grouping can lead to significant reductions of computation load and only minimal performance impact.

The final contribution in this paper will be a method to limit the exposure time to measurement errors. Kalman filter estimates are comprised of every measurement that has been provided to it, so if there was a large error at any point in the past, that error will still influence the current estimate. This paper presents a method to limit the exposure time of measurements occurring farther in the past than a certain time.

We explore these concepts using data GNSS and IMU data collected from an automobile in suburban and highway environments in Calgary, Alberta.

Estimator Design

The PPP algorithm with solution separation has been developed using banks of extended Kalman filters that use dual frequency code phase, carrier phase, and doppler measurements. Additional details regarding the specific implementation can be found in [3], but the implementation is largely a standard PPP filter. The models used for the dual frequency code and carrier phase measurements are as follows:

Dual frequency carrier phase:

$$\Phi_{if}^{(i)} = \left\|x_s^{(i)} - \hat{x}_{rx}\right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} + b_{pwu}^{(i)} - \hat{A}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}_{brdc}^{(i)} + \epsilon^{(i)} \quad (1)$$

Dual frequency code phase:

$$\rho_{if}^{(i)} = \left\|x_s^{(i)} - \hat{x}_{rx}\right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} - \widehat{DCB}_{rx}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}_{brdc}^{(i)} + \epsilon^{(i)} \quad (2)$$

Where

$x_s^{(i)}$- satellite position provided by external precise orbit product

$\hat{x}_{rx}$- estimated receiver position

$\hat{b}_{rx,c}$- estimated receiver clock bias

$b_s^{(i)}$- satellite clock offset provided by external precise orbit product

$m^{(i)}$- tropospheric mapping function

$\widehat{\Delta T}^{(i)}$- estimated delta tropospheric delay

$b_{pwu}^{(i)}$- carrier phase wind-up

$\hat{A}^{(i)}$- estimated float carrier phase ambiguity

$\widehat{M}^{(i)}$- estimated multipath delay on the signal

$\widehat{DCB}_{rx}^{(i)}$- estimated receiver differential code bias per signal (shared across SVs)

$\mathrm{I}^{(\mathrm{i})}$- ionospheric delay/advance

$\mathrm{R_m}$- Other modeled effects. This includes relativistic effects, solid earth tide modeling, satellite antenna phase center offset and variation, ocean loading, modeled tropospheric delay, and any other desired range models. These are strictly modeled and not estimated.

$\hat{\epsilon}_{brdc}^{(i)}$- error due to broadcast navigation message orbit and clock

$\epsilon^{(i)}$- other unaccounted for errors

The estimated states are indicated by a caret over the symbols. Here, the estimated states include the position, velocity, receiver clock biases, tropospheric delay, carrier phase ambiguities, multipath error, receiver differential code bias, and broadcast orbit and clock error. A key element of this system developed for solution separation is that the $\mathrm{R_m}$ term, which is computationally expensive, is computed once using the all-in-view (AIV) solution as an input. The values produced from those operations can then be appropriately applied to the measurements of the subset solutions.

A few states of interest are the differential code biases associated with the receiver. Varying hardware delays within the receiver lead to measurement offsets between the various GPS signals and frequencies. L1 C/A and L2 P semi-codeless do not necessarily follow the same RF paths in the receiver and thus experience different delays. If the only measurements used are dual frequency L1 C/A -L2P, then the hardware delays are combined and pulled into the clock bias estimate. However, if additional signals or single frequency signals are used, then those additional delays must be estimated in order not to introduce error into the position solution. This problem is more extreme when using GLONASS, which uses FDMA signals and has different delays per satellite on the same signal. The estimator implemented in this study estimates separate DCBs for each GPS signal and each GLONASS satellite and signal. The differential delays can be on the order of one meter, so these are important to characterize. For this study, this estimation leads to slower convergence, but in the future, when the receiver can be characterized better before producing protection levels, these values can be calibrated and not estimated. This would lead to even faster convergence.

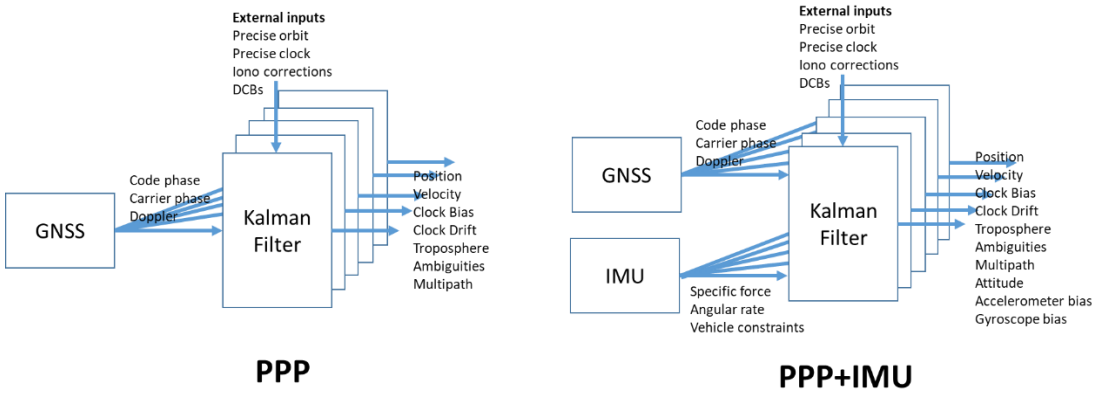

Figure 1: PPP vs PPP+IMU Estimation

Simple checks on the post measurement update residuals are used to exclude measurements that appear to be outliers. The checks and the actual update step is performed iteratively, where the first iteration checks for outliers but can only exclude large outliers. If there are no outliers, then the measurement update can be used as-is. If there are large outliers, those will be removed first, and the measurement update will be performed again. If there are still small outliers, those will then be removed. The reason for the iterative measurement update and multiple thresholds is that the presence of large measurement error, for example a 1-meter carrier phase measurement error caused by a cycle slip that went undetected, can cause excessive residuals on the otherwise good measurements. A single residual check would remove those measurements unnecessarily. This "cleaning" of the measurements that are used by the filter is important in order to keep the probability of fault to an acceptable value.

All of the scenarios in this paper are post-processed, and the estimators have been implemented in MATLAB. The precise orbit and clock data and satellite differential code bias is provided by the IGS MGEX analysis centers [5].

Inclusion of the IMU

Tight coupling of IMU and GNSS has been well studied [6,7,8], so only a brief description relaying specifics of this implementation will follow in this section.

The addition of the IMU essentially replaces the position, velocity, and attitude propagation in the time update step of the Kalman filter as well as adds several new states and measurements in the measurement update step. The filter is an error-state filter implemented in the Earth Centered Earth Fixed (ECEF) frame. Euler angles are used to represent the attitude in ECEF [8].

The measurement update now helps to update IMU bias terms, where the IMU model used is very simple and only estimates accelerometer and gyroscope bias. No scale factor or cross correlation terms are estimated. Because this is not a perfect model of the IMU, additional process noise must be added to the bias terms in order to capture the unmodeled effects. The attitude of the IMU, and because we consider the vehicle to be a rigid body, the vehicle, is estimated. An important assumption made in this study is that the IMU is considered to be fault free, and all of the subset solutions use the same IMU measurements.

Several additional pseudo-measurements are included with the IMU that significantly limit error growth by exploiting the physical nature of the automobile. The first constraint, implemented as a measurement, is the no-slip condition. Essentially, the velocity of the vehicle tires in the body frame cross-track direction must be zero. Non-zero velocity would indicate that the car is sliding, which is an extreme case that we are not considering. The second constraint is that the body frame vertical direction must be zero. Deviation from this constraint could arise from significant motion from the suspension of the vehicle or vehicle separation from the ground. Both of these constraints are considered to be fault-free for the purpose of this study, but this must be further studied in the future.

Integrity algorithm design

We use solution separation-based protection levels derived for Kalman filter navigation solutions as described in [3]. The use of solution separation requires that banks of Kalman filters are run in parallel, each one tolerant to a fault or a set of faults. Much care has been taken to ensure that such banks are not overly computationally demanding, where the main step taken to reduce computational load is the sharing of the computation of modeling across all of the subsets. That is, when the orbit and clock are propagated to the desired time, this is only done once using the all-in-view solution as input, and the orbit and clock state are shared across all subsets. This can be done because in general, these models are not sensitive to the meter or sub-meter level position differences between the subsets. Given the position estimates and covariances output by the bank of Kalman filters, protection levels are computed using algorithms originally developed for ARAIM [3,9,10] which have been since modified for sequential filters.

Ultimately, the protection levels are computed using:

$$PL_{SS,approx2} = \max\left(T_i + Q^{-1}\left(\frac{PHMI}{N * P(H_i)} \right) \sigma^{(i)} \right) \quad (3)$$

Where

Q is the complement of the normal CDF, and Q^{-1} its inverse.

$\sigma^{(i)2}$ is the variance of the estimation error of the subset filter i (the index i=0 corresponds to the all-in-view)

H_i is the fault hypothesis *i*

The threshold T_i is set to meet a predefined probability of false alert P_{fa} under nominal conditions:

$$T_i = Q^{-1}\left(\alpha_i P_{fa}\right)\sigma_{ss}^{(i)} \tag{4}$$

Where

$\sigma_{ss}^{(i)}$ is the standard deviation of the solution separation under nominal conditions.

α_i is the allocation of the probability of false alert to the fault mode. The sum over all modes must not exceed one.

The use of such algorithms requires careful characterization of the environment and the measurements used so that the nominal covariance faithfully reflects the actual error and the fault rates assumed similarly reflect the observed fault rates. The narrow fault rate assumed so far is 10^{-5}, which means that only one-out subsets need be considered.

RESULTS

Both datasets come from the same receiver and IMU setup mounted on an automobile traveling through Calgary, Alberta, Canada. The receiver is a NovAtel OEM 7500 using GPS (L1 C/A – L2P semi-codeless) and GLONASS (L1 C/A- L2P) at 1 Hz. Code phase, carrier phase, and Doppler measurements are used. The IMU is a tactical grade IMU operating at 100 Hz. The lever arm offset between the GNSS antenna phase center and the IMU has been surveyed.

Suburban scenario

The first dataset comes from a suburban scenario. The environment is generally fairly benign and close to open sky, but there are several moments during the collection where partial or full GNSS measurement outages are experienced. Figure 2 shows one of the streets driven in the one-hour data collection. Figure 3 shows the speed of the vehicle throughout the collection, which is generally fairly slow and remains mostly in the 10 to 30 MPH range.

Figure 2: Suburban scenario environment

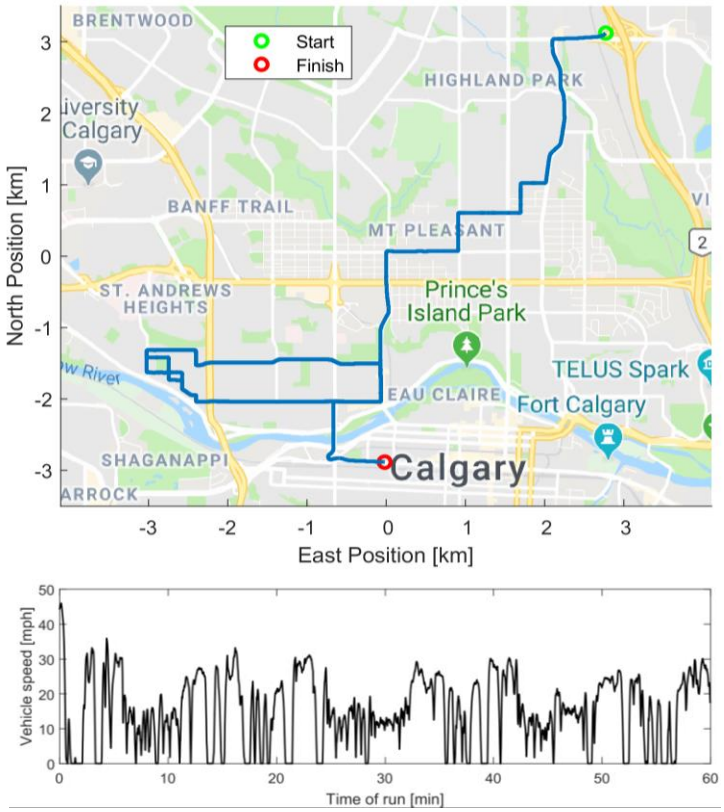

Figure 3: Suburban Calgary scenario path and speed

Figure 4 shows the position error and protection levels produced for the GNSS-only solution in the suburban scenario. The protection level, driven by the filter covariances, converges over the first 10 or so minutes to the 1-2 meter range in the East and North directions and to approximately 4 meters in the vertical direction. These protection levels bound the error quite well. However, one feature that stands out from this figure is that the protection levels jump and must reconverge several times throughout the single hour. This is driven by the partial or full GNSS measurement outages as the vehicle passes under tree cover or next to tall buildings. As signals are occluded and measurements are missing, the covariance swells. The increase in the position covariance is, as is typical for an EKF, driven by the process noise added to the position and velocity states. Without a particular dynamics model or high update rate, the process noise added must be able to account for any dynamics the automobile may go through during the interval between a measurement update and the subsequent time update. When the signals are available later, the carrier phase ambiguities must be re-estimated, which takes time to converge, leading to the characteristic PPP convergence multiple times throughout the data collection.

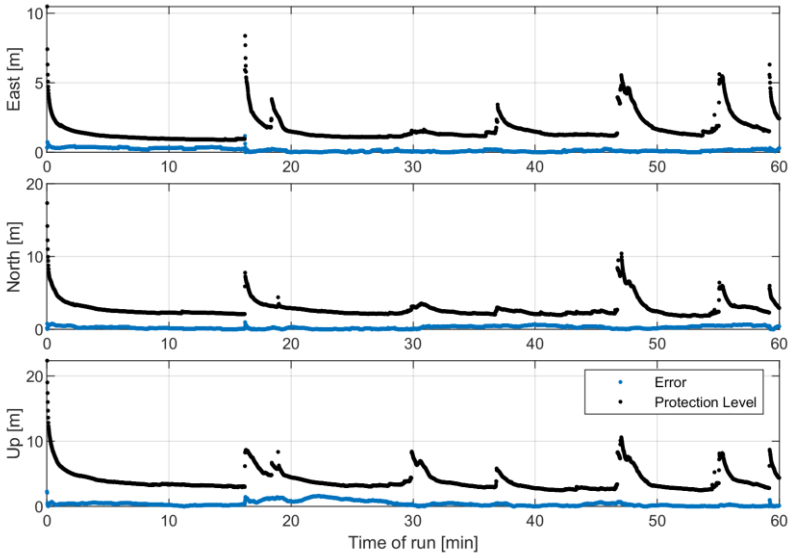

Figure 4: Suburban Scenario Protection Levels and Position Error: PPP only

The inclusion of a tightly-coupled IMU to the PPP solution allows for the short gaps in measurements to be bridged and the protection levels to be kept down, as shown in Figure 5. As before, the protection level must converge upon initialization, but the protection levels do not exhibit the same jumps as they did in the GNSS-only case. The IMU and the land vehicle constraints are able to limit the growth of the position and velocity covariance over the short measurement outages enough to almost entirely smooth out the protection levels.

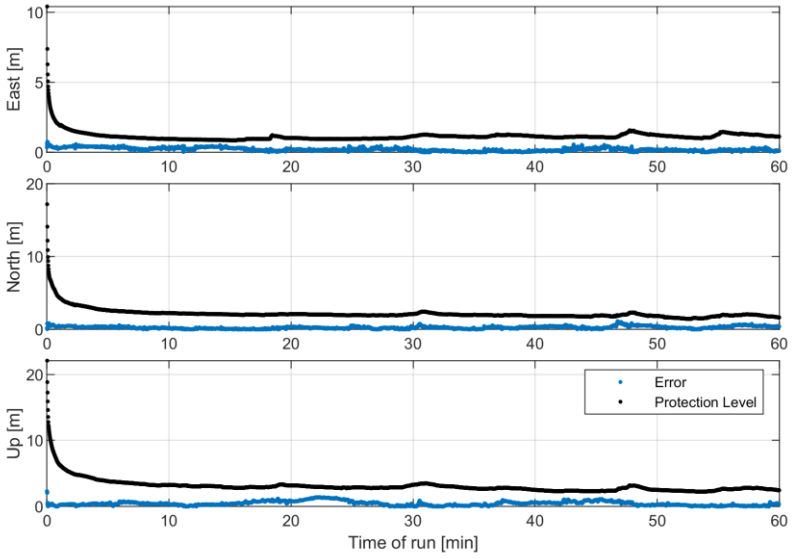

Figure 5: Suburban Scenario Protection Levels and Position Error: PPP+IMU

This growth suppression is, of course, accompanied by an improvement in solution accuracy, especially just following the outages. The protection levels converge to the 1-2 meter range in the horizontal directions and maintain that through the one hour duration. There are sections when the protection level does increase somewhat, and this is driven by worsened GNSS geometry, which does still drive the protection levels long term.

Highway scenario

The next dataset comes from a highway scenario. While the image in Figure 6 looks initially to be generally open sky, the large overhead traffic signs and overpasses in fact cause partial or full measurement outages quite frequently. This environment is much harsher than the suburban scenario because of these features and the frequency with which they are encountered due to the speed of the vehicle. This is also a one-hour data collection. Figure 7 shows that the typical speed throughout this hour is approximately 60 MPH.

Figure 6: Highway scenario environment

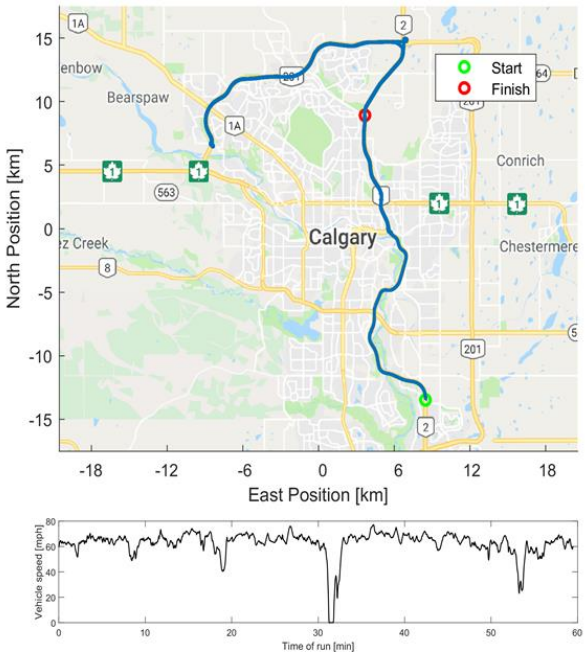

Figure 7: Highway scenario path and speed

Figure 8 clearly shows the effects that the various overhead obstacles have on the ability for the filters to converge. Approximately once per minute, enough measurements are occluded that the covariances, and thus the protection levels, experience a jump. In the East and North directions, the measurement continuity is so poor the protection levels never reach two meters. In the vertical direction, the protection level never reached five meters and is more typically approximately 10 meters.

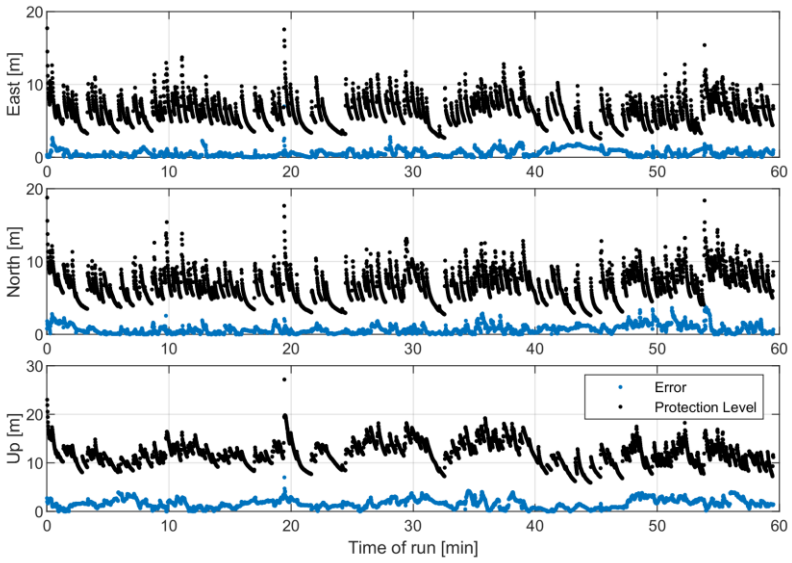

Figure 8: Highway Scenario Protection Levels and Position Error: PPP only

As before, the impact of the addition of the IMU to the state propagation is clear in Figure 9. The protection levels do not experience consistent spikes, and the accuracy is greatly improved. However, due to poorer overall geometry, the protection levels are slightly larger than in the suburban case, typically in the 3-4 meter range in the East and North directions. Still, this is a dramatic improvement when compared to the GNSS-only solution.

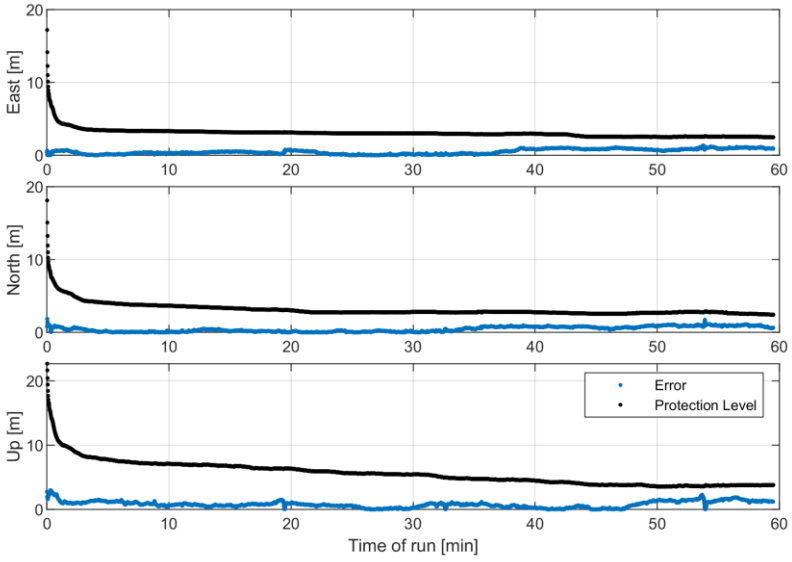

Figure 9: Highway Scenario Protection Levels and Position Error: PPP+IMU

SUBSET GROUPING

The previous section showed that the inclusion of an IMU to the high integrity PPP solution can dramatically improve performance when geometries deteriorate. This section examines grouping of multiple subsets into a single subset to reduce the computational load of the system, where the inclusion of the IMU offsets the performance hit of the grouping. Figure 10 illustrates the concept of grouping subsets. In each of the diagrams, each row is a different filter, and each column represents a different satellite. Filled-in circles indicate that a measurement is used, and empty circles indicate that a measurement is excluded. On the left, each subset excludes only a single subset. When there are six satellites in view, six subsets are required. On the right, each subset excludes two satellite, so that only three subsets are required, but each of the remaining subset filters has worse geometry and higher probability of fault than in the corresponding subset on the left. The case shown in Figure 10 is purely illustrative, as with only six satellites in view, the subset count is not very large, and the grouped subsets have very poor geometries.

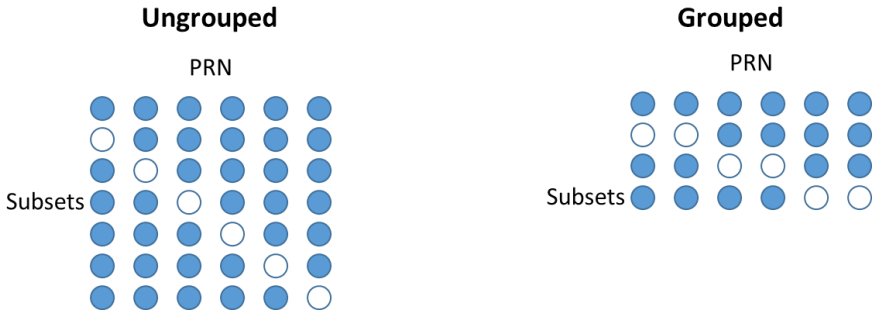

Figure 10: Subset grouping

Nominally, the result of grouping subsets is increased protection levels driven by the increased covariances from the subset filters caused by weakened geometries. When measurements are further occluded or otherwise unavailable, the impacts can become even worse. Figure 11 shows the impact on the protection levels when subsets are grouped into pairs. This data was generated using the suburban scenario, and the non-grouped protection level matches that from Figure 4. Over most of the hour, the impact to the protection levels is very small, but when there are measurement outages, the impact can become more severe. Most notably, 19 minutes into the data collection, there is a measurement outage that particularly affected the protection levels for the paired groupings, leading to protection levels nearly double those of the un-grouped case.

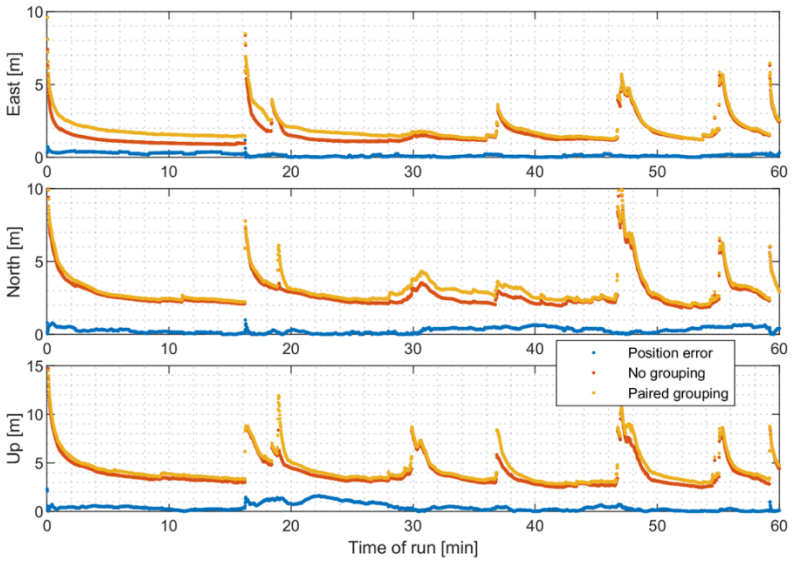

Figure 11: Suburban scenario with grouped subsets- PPP only

The inclusion of the IMU, as shown in Figure 12, mitigates the effects of the poor geometries experienced by some of the grouped subsets. The protection levels produced by grouping pairs are only slightly larger than those of the ungrouped case, but the number of filters that need to be run is halved. This study only considers single satellite fault modes, but if simultaneous faults need to be considered (i.e. two-out subsets), the computational savings become even more significant.

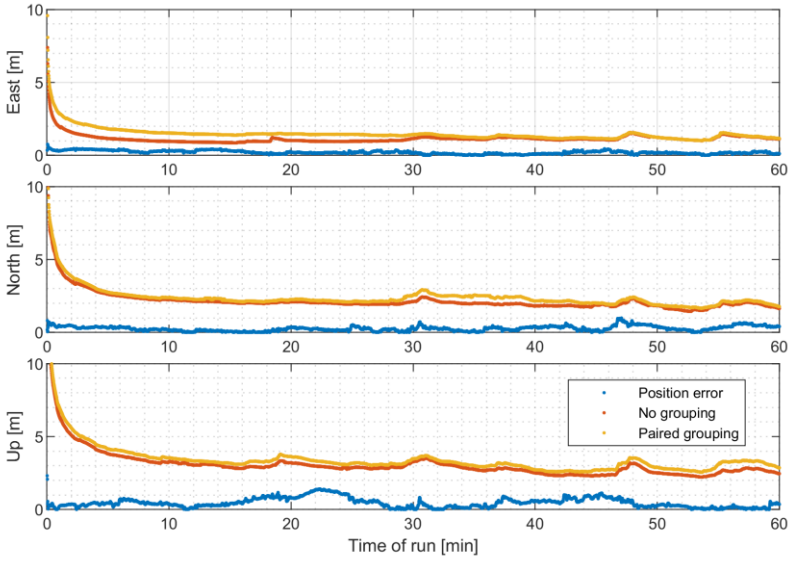

Figure 12: Suburban scenario with grouped subsets- PPP+IMU

REINITIALIZING FILTERS

The final contribution of this paper mainly relates to the removal of older measurements from the filter, most notably measurements from satellites that have set and are no longer available. After a satellite has set, from that time forth, the subset that excludes that satellite will use the same set of measurements as the all-in-view filter, and it might seem reasonable to remove the subset of the set satellite. However, the AIV filter solution is still dependent on the previous measurements from the set satellite, so the subset excluding the set satellite is still required. In order to keep from accumulating subsets indefinitely without removing any, the proposed strategy is to simply reinitialize the AIV and subset filters periodically and run these filters in the background until they have sufficiently converged. When the reinitialized filters have converged, the primary set of filters can be removed, and the reinitialized filters can be promoted to become the primary set of filters.

This process is illustrated in Figure 13. As before, the heavy black lines represent the protection level. The lighter grey lines represent the protection levels of the reinitialized filters in the background. Fifteen minutes into the run, new background filters are spawned. Fifteen minutes after that, when they have sufficiently converged, they become the new primary filters, and the old primary filters are eliminated, and new background filters are initialized. At the 45-minute mark, this process is repeated. By performing this process periodically, old subset filters can be removed, and the measurement exposure time can be limited.

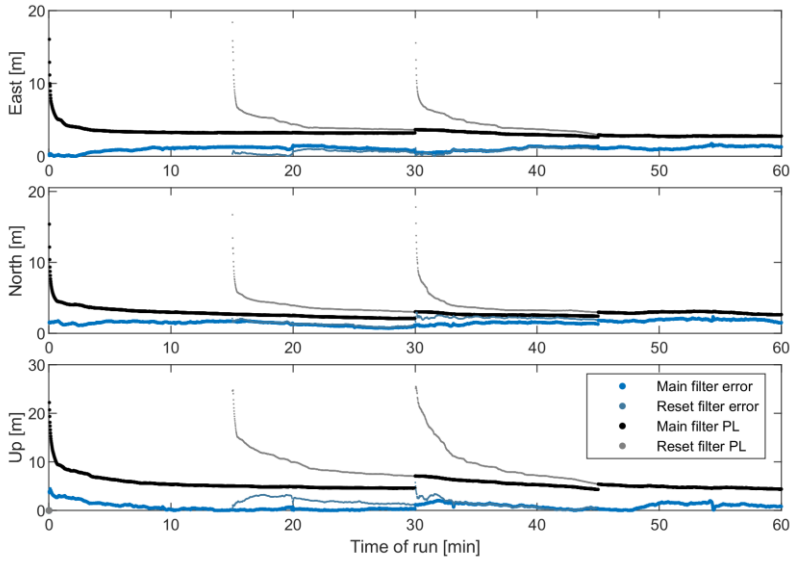

Figure 13: Reinitializing filters

CONCLUSION

A high integrity, high accuracy navigation system using GNSS PPP and tactical grade IMU has been developed and demonstrated. The inclusion of the IMU was shown to improve the continuity and availability of low protection levels when compared to the GNSS- only solution, and 1-2 meter protection levels have been produced in the horizontal directions. In order to reduce the computational load of the solution separation system, the grouping of fault modes (subsets) is also used, which, with the IMU only

leads to a small performance impact to the protection levels. Finally, a method for eliminating subsets that exclude satellites that have set has been proposed.

REFERENCES

[1] Madrid, P. F. Navarro, Fernández, L. Martínez, López, M. Alonso, Samper, M.D. Laínez, Merino, M.M. Romay, "PPP Integrity for Advanced Applications, Including Field Trials with Galileo, Geodetic and Low-Cost Receivers, and a Preliminary Safety Analysis," *Proceedings of the 29th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2016)*, Portland, Oregon, September 2016, pp. 3332-3354.

[2] D. Calle, E. Carbonell, P. Navarro, P. Roldán, I. Rodríguez, G. Tobías, "Facing the Challenges of PPP: Convergence Time, Integrity and Improved Robustness" *Proceedings of the 31th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2018)*,Miami, Florida, September 2018.

[3] K. Gunning, J. Blanch, T. Walter, L. de Groot, L. Norman, "Hexagon Positioning Intelligence, Canada Design and Evaluation of Integrity Algorithms for PPP in Kinematic Applications" *Proceedings of the 31th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2018)*,Miami, Florida, September 2018.

[4] Blanch, Juan, et al. "Reducing Computational Load in Solution Separation for Kalman Filters and an Application to PPP Integrity." *Institute of Navigation International Technical Meeting*, 31 Jan. 2019.

[5] Kouba, Jan. "A guide to using International GNSS Service (IGS) products." (2009).

[6] Petovello, Mark G. *Real-time integration of a tactical-grade IMU and GPS for high-accuracy positioning and navigation*. Calgary, AB: University of Calgary, Department of Geomatics Engineering, 2003.

[7] Jekeli, Christopher. *Inertial navigation systems with geodetic applications*. Walter de Gruyter, 2012.

[8] Groves, Paul D. *Principles of GNSS, inertial, and multisensor integrated navigation systems*. Artech house, 2013.

[9] Blanch, J., Walter, T., Enge, P., Lee, Y., Pervan, B., Rippl, M., Spletter, A., Kropp, V., "Baseline Advanced RAIM User A6gorithm and Possible Improvements," IEEE Transactions on Aerospace and Electronic Systems, Volume 51, No. 1, January 2015.

[10] Blanch, J., Walter, T., Enge, P. "RAIM with Optimal Integrity and Continuity Allocations Under Multiple failures," IEEE Transactions on Aerospace and Electronic Systems Vol. 46, No. 3, July 2010.