Solution separation-based FD to mitigate the Effects of Local Threats on PPP Integrity

Juan Blanch
Stanford University
Stanford, California
blanch@stanford.edu

Todd Walter
Stanford University
Stanford, California
twalter@stanford.edu

Laura Norman
Hexagon Autonomy&Positioning
Calgary, Canada
Laura.Norman@hexagon.com

Kazuma Gunning
Stanford University
Stanford, California
kgunning@stanford.edu

Lance de Groot
Hexagon Autonomy&Positioning
Calgary, Canada
Lance.deGroot@hexagon.com

***Abstract*— In order to obtain protection levels for PPP in multipath prone environments, we formulate a coarse threat model for urban and suburban environments. Based on this threat model, we determine analytically the limitations a of solution separation approach. In particular, we derive for which range of fault rates and fault lag this approach is likely to be feasible. We then describe a solution separation fault detection algorithm adapted to this threat model. Using data collected in urban and suburban environments, we evaluate the benefits of using up to 3 GNSS constellations and dual frequency in PPP combined with an FD algorithm based on solution separation, and in particular its ability to protect against measurement outliers.**

Keywords—Fault detection, measurement outliers, integrity, multipath

I. INTRODUCTION

Many applications, such as autonomous driving or UAV navigation, can greatly benefit from tight guaranteed absolute position error bounds. GNSS Precise Point Positioning [1], which until recently was mostly used as a tool to obtain centimeter-level accuracy, is emerging as a key tool (or set of tools) to achieve decimeter-level high integrity error bounds [2,3,4].

This paper builds on our previous work [4,5,6] on the provision of integrity to PPP using techniques developed for aviation. More precisely, our PPP prototype uses a fault detection and exclusion algorithm based on solution separation. The principle of solution separation is to run a bank of filters, where each filter is fault tolerant to a fault or set of faults. The fault detection statistic is the difference between each of these solutions and the all-in-view solution. In addition to their optimality properties [7], solution separation algorithms offer a straightforward proof of integrity, and good performance in open sky conditions and with occasional outages - when combined with an IMU [7].

In our previous work, we assumed a very simple threat model: up to one satellite could be faulted (after removing obvious outliers based on the innovation residuals). In our results using real measurements, this assumption did lead to adequate protection levels, even in relatively challenging environments. However, such a threat model might not be sufficient to guarantee the integrity of the error bound in environments where many lines of sight might be experiencing large multipath-induced ranging errors. The problem is compounded by the fact that the estimation filter uses measurements over time, such that past undetected faults can have an effect on the current state estimate.

The three goals in this paper are: to refine the threat model to better account for local threats, to adapt, if necessary, our PPP solution separation engine to this threat model, and to evaluate the resulting performance using data collected in challenging environments (urban and suburban road environments).

II. THREAT MODEL AND NOMINAL ERROR MODEL

In order to quantify the integrity risk of the position solution and its error bound, we need to have a trusted model of the nominal errors, and a model of the potential threats.

The nominal model describes the expected measurements when no faults are present. The faults are, by definition, instances where the nominal model is not valid. Note that we may choose to model common (but bounded) faults by inflating the nominal error model. In that sense, the distinction between nominal and faulted conditions is arbitrary.

The most likely fault modes affecting the integrity of the position solution are caused by:

- large code multipath
- undetected cycle slips
- non-line of sight signals
- satellite clock and ephemeris error not captured by the precise ephemeris products
- residual tropospheric delays exceeding nominal model
- residual ionospheric delays exceeding nominal model

Here, we will be targeting integrity risks on the order of 10-3 to 10-4 per hour, so the faults due to the signal in space and the propagation through the ionosphere and the troposphere (which are on the order of 10-5 per hour) can be neglected for now. Here, we will therefore focus on the effect of outliers due to the environment, which are potentially limited in size (a fact that we will exploit). This is a difference with our previous work, where faults could have arbitrarily large errors.

A. Nominal error model

In the nominal state, the error model is assumed to follow a known error model [6], defined by Tables 1, 2, and 3 (second column). For Table 2, the process noise for the multipath state and the carrier phase error state are modeled as random.

TABLE I. INITIAL STATE UNCERTAINTIES

Parameter	***Initial Sigma (meters)***
Tropospheric Wet Delay	0.1
Float carrier phase ambiguity	100
Code phase multipath	2
Carrier phase error state	0.05
Differential code bias	20
Frequency-Dependent DCB	1

TABLE II. PROCESS NOISE SETTINGS

Parameter	***Initial Sigma (meters)***	***Faulted Process Noise (m)***
Tropospheric Wet Delay	$0.002^2/3600^{1/2}$	$0.002^2/3600^{1/2}$
Carrier phase error state	0.01	1
Code phase multipath	0.2	10
Differential Code Bias	0	0
Frequency-Dependent DCB	0	0

TABLE III. MEASUREMENT NOISE SETTINGS

Parameter	Nominal Sigma (m)	Faulted Sigma
Pseudorange	2	10
Carrier phase	0.03	1
Doppler	0.05	0.5
Differential Code Bias	0	0
Frequency-Dependent DCB	0	0

B. Threat model

This work was funded by Hexagon Autonomy&Positioning.

The faults are defined using a very simple model as follows: we assume that up to all measurements from one satellite (we well refer to these are as a measurement series) can transition from the nominal state to a faulted state at a rate R_{fault}, and that the fault will persist during T_{fault}. (Many of the derivations below are valid assuming that T_{fault} represents a mean value of the fault duration, but we will assume for simplicity that they last no more than T_{fault}). In airborne applications, R_{fault} is in the order of 10^{-5} and T_{fault} in the order of 1 hour. For the application we are considering, we expect that R_{fault} will be much higher, and T_{fault} somewhat lower.

A faulted state means that the measurements follow a much more conservative error model defined the third column in Tables 2 and 3. These values are a very coarse estimate, but they do reflect the potential effect of non-line of sight measurements that we have observed in the data. Compared to our previous approach (where errors could be unbounded), this allows us to provide finite (though large) error bounds on the position instead of infinite error bounds in cases where there is not enough redundancy.

C. Outlier rejection based on measurement residuals

As with many practical Kalman filter implementations, the estimator performs checks on the measurement residuals to exclude outliers. These outliers are removed iteratively to reduce the possibility of large outliers affecting the detection of other outliers (because they can drag the solution significantly). It is important to reduce such false alerts in particular for the carrier phase measurements, where a test trip implies a re-initialization of the ambiguity. The exclusion of outliers will reduce the rate of faults that must be mitigated or protected against by other means. However, given our threat model, it is very difficult to provide an estimate of the effect of the outlier rejection on the fault rate. The amount of mitigation will depend on both the type of faults and on the strength of the geometry. For example, cycle slips (and more generally, step faults) will likely be detected under a strong state solution, but they might not if there is not enough redundancy, or too much measurement noise.

It is for this reason that we must assume that some of the faults will not be detected, and must be mitigated through other means. In our case, we expect the solution separation algorithm to mitigate the remaining threats.

D. Fault probabilities, fault rates, error persistence

In our previous papers, we assumed a fixed probability of fault for each measurement series, and we made the assumption that it was sufficient to monitor faults that affect only one of the series at a time. This approach is well adapted to airborne receivers where the threat model is well understood. It is also a first good guess, for other applications. Here, we re-examine this assumption and consider a model that is slightly more sophisticated.

Here we will assume that each measurement series as a fault rate R_{fault}. The probability that a fault occurs anytime within an interval T preceding the current time is given by:

$$P(\text{fault occurs within interval of length } T) = 1 - e^{-R_{fault}T} \approx R_{fault}T \quad (1)$$

A fault will affect the position solution if it lasts long enough to be included in the position solution. We define:

- T_{fault} the mean length of a fault and
- T_{filter} the period beyond which a measurement will not have an impact on the solution

Without modification, in a Kalman filter the time T will increase indefinitely. As a consequence, within this model, the probability of having a fault is guaranteed to affect so many time series that it will not be possible to obtain a fault free subset with sufficient probability. To prevent this from happening, there are at least two approaches. One consists in using a batch approach instead of a Kalman filter. In this approach, any measurements older than a fixed T are not included in the solution. This way we are guaranteed that any fault occurring to measurements older than T will have no impact on the state estimate. A second approach was described in [6], and it consisted in keeping two estimation filters running in parallel that re-initialize at an interval of length T_{filter}, where the re-initialization is staggered every $T_{filter}/2$. This way, the receiver can always use one of the state estimates that has converged, and the state estimate will not be exposed to faults that have happened before T_{exp}. More details on this approach can be found in [6].

E. State probability

With this model, the probability that a fault in one measurement series i affects the state estimate at a given time is therefore given by:

$$P_{fault,i} = R_{fault,i}\left(T_{fault} + T_{filter}\right) \quad (2)$$

For this paper, we will also assume that the probability of having a fault in one satellite measurement series is independent of the others. As a consequence, the probability that two measurement series are affected is given by:

$$\text{Prob}(\text{fault in } i \text{ and } j \text{ measurement series}) = P_{fault,i}P_{fault,j} \quad (3)$$

This formula can be generalized to an arbitrary number of measurement series.

F. Probability that multiple faults affect the state estimate over a finite period

The previous calculation describes the probability that a state estimate at a given time is affected by a fault, and this will be useful later when we evaluate the integrity risk contribution of the fault modes that are actively monitored.

However, we also need an upper bound of the contribution of fault modes that are not monitored over a given exposure time. We could derive this probability from the state probability and the required time to alert, but that turns out to be exceedingly conservative.

Instead, to compute such a bound, we consider the probability of having q simultaneous faults at any point during a given exposure time T_{exp}. To go further, we compute the rate of such faults. The rate of a multiple fault is given by (see Appendix B):

$$R_{composite} = \prod_{k_1,\cdots,k_q} P_{fault,k_i} \sum_{r=1}^{q} \frac{1}{T_{fault,k_i} + T_{filter}} \quad (4)$$

The average length of such a fault (as it affects the filter) is given by:

$$T_{composite} = \left(\sum_{r=1}^{q} \frac{1}{T_{fault,k_i} + T_{filter}} \right)^{-1} \quad (5)$$

The probability that the fault affects the state estimate in the interval $[0\ T_{exp}]$ is the probability that it appears anywhere in the interval $[-T_{composite}\ T_{exp}]$. Therefore, the probability of the fault impacting the exposure interval is given by:

$$P_{fault,composite} = \prod_{k_1,\cdots,k_q} P_{fault,k_i} \left(1 + \sum_{r=1}^{q} \frac{T_{exp}}{T_{fault,k_i} + T_{filter}} \right) \quad (6)$$

For the case where all the faults have the same rate and duration, we have:

$$P_{fault,composite} = P_{fault}{}^{q} \left(1 + q \frac{T_{exp}}{T_{fault} + T_{filter}} \right) \quad (7)$$

Now let us consider all such composite faults. There are $\binom{n}{q}$ of them, where n is the number of satellite measurement series. The probability of having q or more simultaneously affecting our state estimate is given by:

$$\begin{aligned} &P(q \text{ simultaneous faults affecting state estimate}) \\ &= \binom{n}{q} P_{fault}{}^{q} \left(1 + q \frac{T_{exp}}{T_{fault} + T_{filter}} \right) \\ &= \binom{n}{q} R_{fault}{}^{q} \left(T_{fault} + T_{filter} \right)^{q} \left(1 + q \frac{T_{exp}}{T_{fault} + T_{filter}} \right) \end{aligned} \quad (8)$$

A fault detection algorithm monitoring up to *q-1* simultaneously faulted satellites (in the sense that they have been affected in the past $T_{fault}+T_{filter}$) would not necessarily detect faults affecting q or more satellites. For this reason, the above formula represents an upper bound of the integrity risk contribution from the modes that are not monitored.

For example, let us consider a situation where:

$$R_{fault,i} = 10^{-2} / hour$$

$T_{fault} + T_{filter} = 120s$

$T_{\exp} = 1$ hour

n = satellites

The probability of having two simultaneous faults (q=2) impacting the state estimate over the exposure time T_{exp} is, using Equation (8), $3.0*10^{-4}$.This means that if we have an integrity target of 10^{-3}, it would be sufficient to monitor the single faults. For an integrity target of 10^{-4}, it would not be sufficient.

In Figure 1 we plot this probability with a fixed $T_{fault}+T_{filter}$=120s and n=10 for a range of fault rate values and number of simultaneous faults. In addition, we plot a notional threshold value of 5 x 10^{-5}, below which we can assume that the faults won't need to be monitored. For a given rate, only the values of *q* with a value below the threshold will lead to an algorithm with the potential of meeting the integrity requirement.

Similarly, in Figure 2 we keep a constant rate of 0.1/hour and vary the fault persistence. Each of the lines corresponds to *q*, the number of simultaneous faults.

From these few examples it becomes clear that guaranteeing low integrity risk values will require the active monitoring of many fault modes, even for moderate fault rates. For large fault rates and long fault persistence times, this approach appears to be impractical. For example, from Figure 2, we can see that a rate of 10^{-1}/hour and $T_{fault}+T_{filter}$=15 minutes would require the active monitoring of all modes composed of up to 4 simultaneous faults. For our implementation, we assumed a rate of 10^{-2}/hour, which means that for a target 10^{-4} per hour, at least two-out subsets need to be monitored.

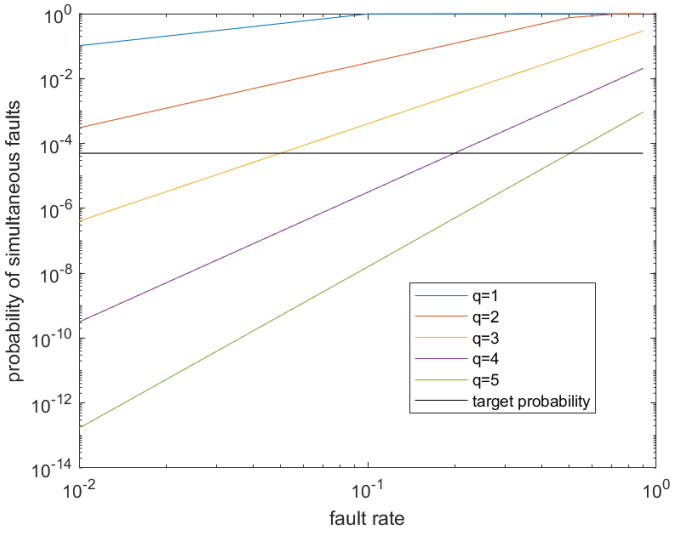

Fig. 1. Probability of q simultaneous faults over an interval of 1 hour for 10 satellites as a function of the fault rate per hour.

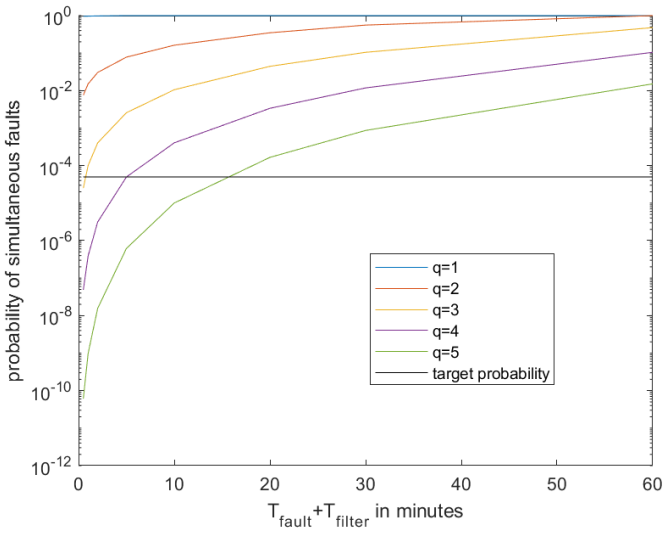

Fig. 2. Probability of *q* simultaneous faults over an interval of 1 hour for 10 satellites as a function of the fault persistence.

G. Protection Level Equation

The previous subsection has determined the depth of subsets that need to be monitored. A set of protection level equations was described in [4]. Techniques to lower the computational load were described in [5]. For this paper we chose the simplest form, which is given by:

$$PL_{SS,approx2} = \max\left(T_i + Q^{-1}\left(\frac{PHMI}{N * P_{fault}}\right)\sigma^{(i)}\right) \quad (9)$$

$$T = Q^{-1}\left(\frac{P_{fa}}{N}\right)\sigma_{ss}^{(i)}$$

where:

$\sigma^{(i)}$ is the standard deviation of the position error in the coordinate of interest

N is the number of subsets

PHMI is the integrity risk per sample (after removing the contribution of the modes that are not monitored). We will assume that there are 3600 samples in one hour, such that PHMI = 10^{-4}/3600 = 2.7 10^{-8}.

P_{fa} is the false alert per sample. If we target a probability of false alert per hour of 10^{-3}, then per sample we have Pfa =10^{-3}/3600 = 2.7 10^{-7}.

For the implementation shown in the next section we have:

$$P_{fault} = R_{fault}\left(T_{fault} + T_{filter}\right) = 10^{-2} * \frac{120}{3600} \simeq 3.3 \cdot 10^{-4}$$

All these values are at this point notional. They only represent a first estimate, both for the requirements and the fault rates.

III. Implementation in a Multi-Constellation PPP Prototype

A complete description of our prototype has been described before [4], [5], [6]. 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 the algorithm developed for solution separation is that the term including relativistic effects, solid tides (among other terms), which is computationally expensive, is computed once using the all-in-view (AIV) solution as an input. Precise clocks and orbits were from the Center for Orbit Determination in Europe (CODE) for all constellations (IGS MGEX product). We used the IGS satellite antenna phase center offsets and variations. Ionospheric TEC maps were also from CODE.

We consider two GNSS data sets, one representing an urban environment and a suburban one. The data sets are described in Table 4.

TABLE IV. Data Sets Summary

	Urban	Suburban
Receiver	NovAtel OEM 7500	Trimble BX 940
Amount and date	20 minutes of driving data on March 1, 2018	15 minutes of driving data in July 2019
Signals	GPS (L1 C/A -L2P semi-codeless), GLONASS (L1 C/A-L2P) at 1 Hz	GPS (L1 C/A -L2P semi-codeless, L5), GLONASS (L1 C/A-L2P) at 1 Hz, Galileo (E1/E5)
Truth	NovAtel OEM729 with tactical-grade IMU with forward and reverse processing	Provided by the Natural Resources Canada Canadian Spatial Reference System Precise Point Positioning (CSRS-PPP) service

A. *Results for the suburban environment*

This data was collected in roads with heavy foliage, such that the masking angle was above 25 degrees. However, we did use three constellations, which results in about 10 satellites in view most of the time. We did not have a truth reference for all time steps (the IMU measurements were not integrated in the RTK solution).

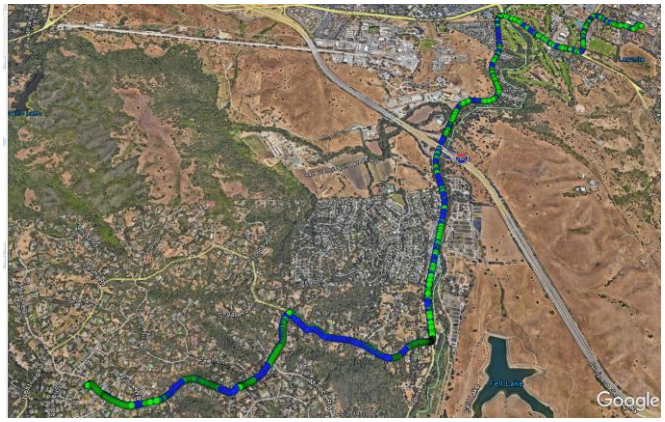

Fig. 3. Path corresponding to the sub-urban scenario

Figure 4 shows three time series corresponding to the dual constellation case (GPS+GLONASS):

- The position error (blue)
- The protection level derived from the nominal model only (red)
- The protection level computed using the solution separation algorithm described above

We can see that, although the nominal error bound covers the position errors, the margin is sometimes thin. In contrast, the protection level appears to offer a larger margin. However, the protection levels are rather large. Figure 5 shows the same scenario using Galileo in addition to GPS and GLONASS. The protection levels are still high due to the large number of cycle slips, but they have been significantly reduced.

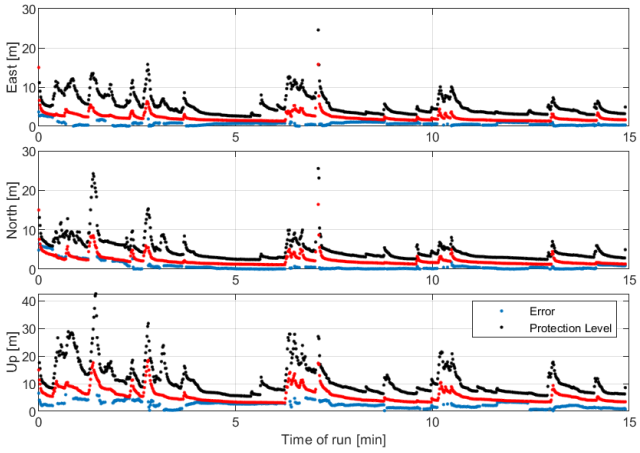

Fig. 4. Position errors, nominal error bounds (red), and protection levels with two constellations (GPS and GLONASS)

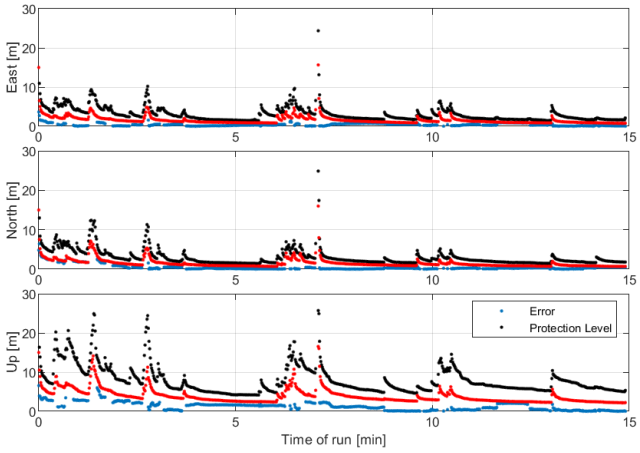

Fig. 5. Position errors, nominal error bounds (red), and protection levels with three constellations (GPS and GLONASS)

B. *Results for the urban scenario*

Figure 6 shows the path followed by the vehicle in the urban scenario and Figure 7, a summary of the measurements collected and used by the filter in the urban environment. This represents an extremely challenging GNSS environment, where the cycle slips are constantly occurring and the number of measurements drop to zero in several occasions. It is therefore clear that the measurements do not follow the nominal model at the assumed rate. Still, we wanted to test whether our approach helps determine the trust in each of the position solutions.

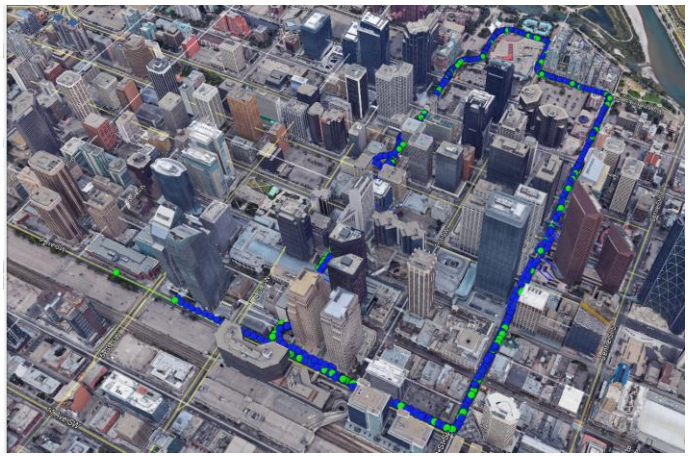

Fig. 6. Path corresponding to the urban scenario.

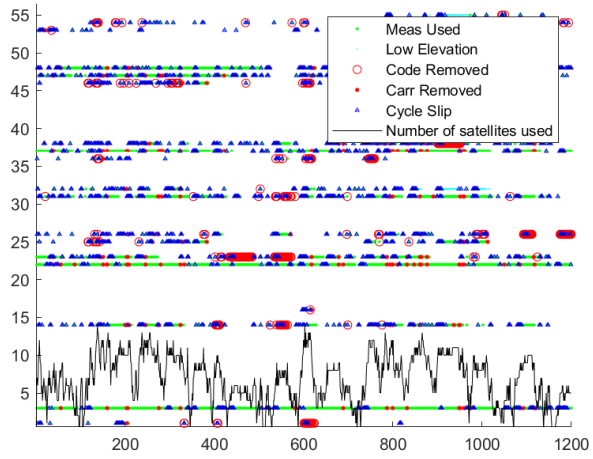

Fig. 7. Overview of GNSS measurements for the urban scenario.

Figure 8 shows the position error, the nominal error bound, and the protection level (as with Figures 4 and 5 in the suburban scenario).

We can see that, in many instances, the error exceeds the protection level derived from the nominal model only (close to 10% of the time). In contrast, the protection level derived from the solution separation covers the observed errors. As pointed out, the threats affecting the solution are much worse than what is assumed by the algorithm and we therefore would not expect it to meet the desired integrity figures, but it appears to add some protection to the user.

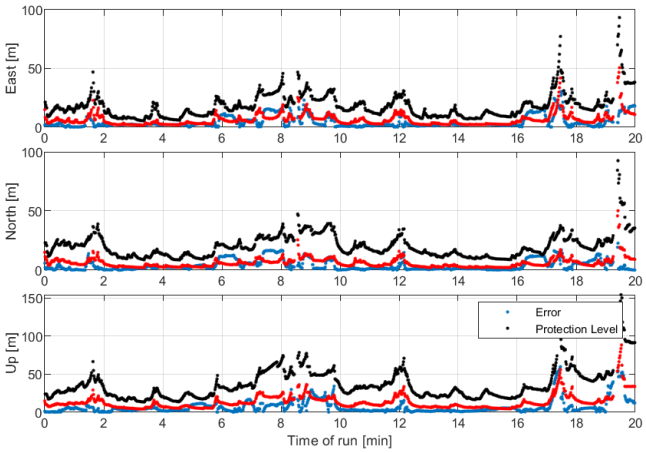

Fig. 8. Position errors, nominal error bounds (red), and protection levels.

For the final version of this paper, we plan to process a larger number of data sets.

IV. Summary

In this paper we examine whether multi-constellation GNSS combined with fault detection techniques can help mitigate the effects of multipath and non-line of sight measurements on precise point positioning. In order to obtain protection levels for PPP in multipath prone environments, we develop a coarse threat model for urban and suburban environments. We then design a solution separation FDE algorithm adapted to this threat model.

We briefly evaluate the benefits of using up to 3 GNSS constellations and dual frequency in PPP combined with an FDE algorithm based on solution separation, and in particular its ability to mitigate multipath in urban and suburban driving conditions. In the suburban scenario, we find that the nominal model generally covers the position errors, from which we deduce that the algorithm would be tolerant to undetected faults in up to two satellites. We also find that adding a third constellation (Galileo) to the solution strengthens both the accuracy and the protection levels by more than 50%. In the urban scenario, which was very challenging, we find that the protection levels help protect the user against large position errors.

Acknowledgment

This work was funded by Hexagon Positioning Intelligence. We would like to thank Lance De Groot and Laura Norman from NovAtel for providing the urban GNSS data.

References

[1] Kouba, J. and Heroux, P. (2001). Precise Point Positioning Using IGS Orbit and Clock Products, GPS Solutions, vol. 5, no. 2, pp. 12-28.

[2] 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.

[3] 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.

[4] K. Gunning, J. Blanch, T. Walter, L. de Groot, L. Norman, "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.

[5] Blanch, J., Gunning, K., Walter, T., De Groot, L., Norman, L., "Reducing Computational Load in Solution Separation for Kalman Filters and an Application to PPP Integrity," *Proceedings of the 2019 International Technical Meeting of The Institute of Navigation*, Reston, Virginia, January 2019, pp. 720-729.

[6] Gunning, Kazuma, Blanch, Juan, Walter, Todd, de Groot, Lance, Norman, Laura, "Integrity for Tightly Coupled PPP and IMU," *Proceedings of the 32nd International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2019)*, Miami, Florida, September 2019, pp. 3066-3078.

[7] Blanch, Juan, Walter, Todd, Enge, Per, "Theoretical Results on the Optimal Detection Statistics for Autonomous Integrity Monitoring", NAVIGATION, Journal of The Institute of Navigation, Vol. 64, No. 1, Spring 2017, pp. 123-137.

APPENDIX

A. *PPP filter*

The measurements used in the EKF are single and dual frequency code and carrier phase measurements along with single frequency Doppler measurements. Continuity of L1 carrier phase measurements is significantly better than that of L1/L2 measurements. Using L1-only measurements in addition to dual frequency measurements leads to much smoother covariances and protection levels over time that do not jump up upon loss of L2 measurements. The L1 measurements require ionospheric delay estimates. In this case, IGS TEC maps are used for both code and carrier measurements. However, the filter could be further simplified by only using L1 carrier phase measurements, not including an ionospheric estimate, and adding more process noise to the carrier phase error state due to the temporal changes in the ionospheric delay. L1 Doppler measurements are included for a more direct measurement of the receiver velocity. The code and carrier phase measurements are modeled as follows:

Dual frequency carrier phase:

$$\Phi_{if}^{(i)} = \left\| x_s^{(i)} - \hat{x}_{rx} \right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} + b_{pwu}^{(i)} - \hat{A}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}^{(i)} \tag{10}$$

Dual frequency code phase:

$$\rho_{if}^{(i)} = \left\| x_s^{(i)} - \hat{x}_{rx} \right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} - \widehat{DCB}_{rx}^{(i)} - f_j\widehat{FDCB}_{rx}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}^{(i)} \tag{11}$$

Single frequency carrier phase:

$$\Phi^{(i)} = \left\| x_s^{(i)} - \hat{x}_{rx} \right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} + b_{pwu}^{(i)} - \hat{A}^{(i)} - \mathrm{I}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}^{(i)} \tag{12}$$

Single frequency code phase:

$$\rho^{(i)} = \left\| x_s^{(i)} - \hat{x}_{rx} \right\| + c\left(\hat{b}_{rx,c} - b_s^{(i)}\right) + m^{(i)}\widehat{\Delta T}^{(i)} - \widehat{DCB}_{rx}^{(i)} - f_j\widehat{FDCB}_{rx}^{(i)} + \mathrm{I}^{(i)} + \mathrm{R_m} + \widehat{M}^{(i)} + \hat{\epsilon}^{(i)} \tag{13}$$

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)

f_j- GLONASS signal frequency channel number from -7 to 6

$\widehat{FDCB}_{rx}^{(i)}$ - estimated frequency-dependent GLONASS differential code bias per signal (shared across SVs)

$\mathrm{I}^{(i)}$- 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.

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

Receiver differential code bias states have been included to accommodate the inclusion of single frequency measurements. For non-GLONASS constellations, one signal is assigned as the reference, and a constant DCB is estimated for every other signal. There is no process noise added to the DCB state. For GLONASS, a frequency-dependent DCB is included for each signal to account for local delays that are a function of the satellite frequency channel.

B. *Fault rate of composite faults*

We demonstrate the formula for two faults. The general formula can be deduced by induction. Let us consider two events 1 and 2, with rates R_1 and R_2 and mean time to detect m_1 and m_2. In order to determine the rate of the combined fault, let us consider the probability that the fault appears in an interval Δt.

We define P_i as the probability that event i is present at a given time. We have the relationship:

$$P_i = R_i m_i \tag{14}$$

There can be three mechanisms for the composite fault to appear:

- fault 1 was already present and 2 starts in the infinitesimal interval Δt, or
- fault 2 was already present and 1 starts in the infinitesimal interval Δt
- both faults 1 or 2 appear in the interval Δt

The probability of 1 and 2 appearing in the interval Δt is therefore:

$$\begin{aligned} &P\left(\text{event 1\&2 appears in } [0,\Delta t]\right) \\ &= P_1 R_2 \Delta t + P_2 R_1 \Delta t + \left(R_1 \Delta t\right)\left(R_2 \Delta t\right) \end{aligned} \tag{15}$$

To obtain the rate, we divide this expression by Δt and make it go to 0:

$$R_{12} = \lim_{\Delta t \to 0} \frac{P_1 R_2 \Delta t + P_2 R_1 \Delta t + \left(R_1 \Delta t\right)\left(R_2 \Delta t\right)}{\Delta t} = P_1 R_2 + P_2 R_1 \tag{16}$$

Using Equation (14), we can write:

$$\begin{aligned} R_{12} &= P_1 R_2 + P_2 R_1 = P_1 P_2 \frac{1}{m_2} + P_1 P_2 \frac{1}{m_1} \\ &= P_1 P_2 \left(\frac{1}{m_1} + \frac{1}{m_2}\right) \end{aligned} \tag{17}$$

To compute the mean fault duration of the composite fault m_{12}, we consider the probability of having 1 and 2 at a given time. Using Equation (14) again, we have:

$$P_{12} = R_{12} m_{12} \tag{18}$$

Combining Equations (17) and (18), we get:

$$m_{12} = \left(\frac{1}{m_1} + \frac{1}{m_2}\right)^{-1} \tag{19}$$