

A Framework for Integrity Monitoring for Positioning Through Graph-Based SLAM Optimization [†]

Sam Bekkers ^{*} and Heiko Engwerda ^{*} 

Department of Aerospace Systems Intel & Space Capabilities, Royal Netherlands Aerospace Center (NLR),
1059 Amsterdam, The Netherlands

^{*} Correspondence: sam.bekkers@nlr.nl (S.B.); heiko.engwerda@nlr.nl (H.E.)

[†] Presented at the European Navigation Conference 2025 (ENC 2025), Wrocław, Poland, 21–23 May 2025.

Abstract

As satellite navigation systems show vulnerabilities in specific circumstances such as urban canyons or jamming and spoofing situations, additional sensors such as cameras may be incorporated on the platform. Despite advancements in the robotics and computer vision community, which have led to increasingly accurate Simultaneous Localization and Mapping (SLAM) positioning solutions, visual navigation has its own vulnerabilities. It therefore remains of critical importance for many applications to study the integrity of fused navigation algorithms and their components, which is done less for SLAM than for satellite navigation. In this paper, a framework for integrity monitoring (IM) of a visual SLAM algorithm is proposed. A sensor-level IM scheme analyses feature reprojection errors. It is demonstrated that, in dynamic environments, multiple hypotheses can be generated from different subsets of extracted features. Additionally, the factor graph-based framework employs a fusion-level IM scheme which deals with these multiple hypotheses and selects the most probable one by calculating the sum of weighted measurement residuals. These concepts are applied to scenarios from real and simulated experiments in order to demonstrate applicability.

Keywords: factor graph optimization; SLAM; dynamic environments; robustness; integrity monitoring

1. Introduction

Ensuring safe operations of platforms such as Unmanned Aerial Vehicles (UAVs) relies heavily on the availability of accurate and reliable navigation and localization capabilities of the platform. Although the Global Navigation Satellite System (GNSS) can provide continuous positioning information to these platforms in nominal conditions, the integrity of the GNSS may not be guaranteed in urban canyons where multipath or Non-Line-of-Sight effects occur, or areas contaminated with GNSS signal interference. Various integrity monitoring (IM) algorithms have been developed to indicate the health and quality of the estimated position solution and give timely warnings when the solution quality degrades. Receiver Autonomous Integrity Monitoring (RAIM), with its two variants, Solution-separation RAIM and Residual-based RAIM [1], relies on measurement redundancy to provide a sense of the level of agreement between measurements. Autonomous Integrity Monitoring Extrapolation (AIME) [2] specifically detects ramp failures using a sliding window, and an improved AIME method [3] varies the sliding window length to overcome the issue of balancing the detection sensitivity and response time. Solutions



Academic Editor: Tomasz Hadaś

Published: 25 February 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

based on a Factor Graph Optimization (FGO) improve measurement redundancy by taking historical measurements into account [4]. Proceeding from this idea, ref. [5] designed a Fault Detection and Exclusion (FDE) scheme for pure GNSS navigation applied over multiple epochs in a sliding window and demonstrated that the proposed FDE method can help to improve positioning accuracy.

Simultaneously, fusion of multi-sensor data combines complementary advantages of various sensors [6]. In [7], it is demonstrated how integrity monitoring can be conducted for GNSS navigation augmented with inertial measurements. Additionally, visual SLAM algorithms show great potential for navigation, although they rely heavily on the quality of the captured images for accurate motion estimation. For robust feature-based SLAM, the feature set can be split to estimate camera rotation and translation separately [8]. In [9], a heterogeneous factor graph is constructed based on multiple sensor modalities. Using switch variables, a chi-square test variable is satisfied within the optimization, yielding high positioning accuracy.

This work investigates an IM scheme on the sensor level for visual odometry to generate multiple hypotheses in dynamic environments, based on the redundancy in visual feature measurements. A framework is discussed that resolves these multiple hypotheses using a higher-level sensor fusion integrity monitor by performing a χ^2 -like residual test. The method is based on the open-source SLAM library RTAB-map [10] and incorporates GNSS measurements. This increases the global consistency of the SLAM solution and provides the necessary measurement redundancy to the global system-level integrity monitor. The work is presented as a framework, as both of the IM schemes could be replaced by other IM strategies; different strategies for generating multiple hypotheses can be implemented, as well as different chi-square test statistic detectors.

2. Navigation Method

In this paper, a factor graph-based solution is used for the multimodal sensor fusion of GNSS and visual measurements to enable maximum-a-posteriori (MAP) inference of the navigation state parameters. This is a graph-based SLAM method that leverages the sparse nature of the underlying navigation problem, rather than optimizing all observed poses and landmarks like global Bundle Adjustment (BA) does. This allows for maintaining multiple hypotheses in multiple parallel factor graphs to detect and exclude faults in a computationally feasible manner.

2.1. Visual Odometry

The work is built upon the existing RTAB-map library, which performs visual odometry using RGB-D images. By extracting and matching features from consecutive images $k - 1$ and k , the motion between the two camera frames is estimated by minimizing the reprojection error of these features. More specifically, the optimization step employs a Perspective-n-Point (PnP) algorithm from the OpenCV library [11], which is illustrated in Figure 1.

In an iterative manner, it estimates a camera pose given a set of features in 3D world coordinates observed in an RGB-D image $k - 1$, and their matched 2D features in the RGB image k . Using a projection function $\pi(\cdot)$ that projects the 3D feature points onto a proposed 2D camera plane, the reprojection error is minimized using a Levenberg-Marquardt optimizer. This returns a camera rotation and translation R, t with respect to the 3D feature points, which are expressed in the local coordinate frame of image $k - 1$,

and thus provides information about the relative camera motion between those frames. Mathematically, the PnP algorithm is defined as:

$$R^*, t^* = \underset{R, t}{\operatorname{argmin}} \sum_{i=1}^n \| \mathbf{u}_i - \pi(R\mathbf{X}_i + t) \|^2 \quad (1)$$

in which $\mathbf{u}_i$ is the set of observed feature coordinates in 2D image coordinates, n is the number of matched features across image $k - 1$ and k , and $\mathbf{X}_i$ is the set of corresponding features in 3D world coordinates. For robustness towards outliers, a Random Sample Consensus (RANSAC) scheme is implemented in RTAB-map. After every optimization iteration, features with a reprojection error exceeding some predefined threshold are discarded from the measurement set over which optimization is performed. This improves the accuracy of motion estimation by excluding faulty measurements, such as misassociations of features across images or features with an unacceptably high noise. Additionally, it provides some robustness towards dynamic environments, as features corresponding to dynamic obstacles are likely to be correctly excluded since they do not match the estimated motion based on the rest of the feature measurements. However, when a dynamic obstacle occupies a substantial part of the image, RANSAC might select the wrong features for motion estimation. Section 3.1 will elaborate further on the RANSAC scheme, its drawbacks, and the proposed modifications.




Figure 1. The Perspective-n-Point algorithm as implemented by OpenCV. Given correspondences between 3D object points in the world coordinate system (depicted in orange) and 2D image points on the camera plane (depicted in yellow) and known intrinsic camera calibration parameters, the PnP algorithm finds the 6 degree-of-freedom (DoF) pose of the camera in world coordinates.

2.2. Factor Graph Optimization

The factor graph is a graphical probabilistic model that specifies a joint probability density as a product of factors. The factor graph optimization algorithm solves for both current and historical data, increasing its global consistency as compared to a Kalman filter. In a setup where multi-sensor data is fused in a factor graph, the measurement factors are constructed using the sensor observations, either directly when the sensor output concerns position information, like GNSS, or indirectly by constructing motion constraints from the observed measurements as done for Visual Odometry (VO). Maximum Likelihood Estimation (MLE), i.e. finding the most probable position solution given the observed measurements, is done by maximizing the posterior density of the joint density distribution $\phi(X_i)$:

$$X^{MAP} = \underset{X}{\operatorname{argmax}} \prod_i \phi_i(X_i) \quad (2)$$

Assuming that the noise models of the sensors follow a Gaussian distribution under nominal conditions, finding the MLE turns to minimizing the following non-linear least squares problem:

$$X^{MAP} = \underset{X}{\operatorname{argmin}} \sum_i \|h_i(X_i) - z_i\|_{\Sigma_i}^2 \quad (3)$$

in which $h_i(X_i)$ and z_i are the estimated and measured values of the factor parameters, the platform position in this paper. Based on this theoretical framework, the multi-sensor factor graph architecture is constructed as depicted in Figure 2. This factor graph is incrementally extended by incorporating GNSS- and VO-factor motion estimates. A prior factor describes the relationship between the nodes and the global coordinate frame, and loop closure (LC) constraints represent the place recognition constraints added when the platform recognizes a previously visited location.




Figure 2. Schematic overview of constructed factor graph. Factors are indicated by blue boxes, the factor type is mentioned above each factor.

The output of the visual odometry module, R and t , is used to construct a homogeneous transform ${}^G T_n$, which represents the pose estimated by VO of the camera at step n in the global frame G . Together with the previous pose ${}^G T_{n-1}$, it is used to formulate the f^{VO} factor, which provides a 6 DoF motion estimate :

$$z_n^{VO}(x_n, x_{n-1}) = {}^G T_n \times ({}^G T_{n-1})^{-1} \quad (4)$$

$$f^{VO}(x_n) \propto \exp\left(\left\|h_n^{VO}(x_n) - z_n^{VO}\right\|_{\Sigma_{VO}}^2\right) \quad (5)$$

The GNSS factor provides an absolute position measurement to the user, by providing a 3 DoF constraint, as it does not constrain the orientation of a pose. The GNSS factor is modeled as:

$$f^{GNSS}(x_n) \propto \exp\left(\left\|h_n^{GNSS}(x_n) - z_n^{GNSS}\right\|_{\Sigma_{GNSS}}^2\right) \quad (6)$$

3. Integrity Monitoring

This work implements integrity monitoring (IM) schemes on two levels. A sensor-level IM algorithm focuses on individual VO measurements based on the feature measurement redundancy. It investigates the reprojection errors and residual distributions and forms the basis of a multiple-hypothesis IM scheme which investigates the integrity of a single visual odometry measurement. A fusion-level IM algorithm validates to what extent these VO constraints agree with other factors in the factor graph, which could originate from a different sensor modality or a loop closure constraint. Figure 3 illustrates the general navigation process flow and at which stage in the process the two IM algorithms are applied.




Figure 3. Complete SLAM process flow with indicated integrity monitoring systems on local (sensor) and global (sensor fusion) level. Light blue boxes indicate potential future extensions of the framework presented in this work.

3.1. Visual Odometry Integrity Monitoring

A more detailed scheme of the sensor-level IM monitor is depicted in Figure 4. As mentioned before, the VO module performs extraction and matching of visual features across consecutive images. Noise or faults in these features' measurement may occur due to various reasons, such as image blur due to motion, misassociation of features across images during the matching phase, or the influence of dynamic obstacles in a scene. Figure 5 illustrates how dynamic obstacles in a scene potentially confuse the RANSAC scheme. Selected inliers are drawn in green; the red and yellow points are excluded from the PnP algorithm.




Figure 4. Flow of the local VO integrity monitoring algorithm. Features are indicated by the stars. In the first epoch, green stars illustrate inliers.




Figure 5. (a) Shows a frame in which RANSAC mistakes the features belonging to the cyclist for inliers (green) and removes the static features belonging to the pavement. (b) Shows the same cyclist a couple of frames later, where RANSAC has correctly segmented the dynamic features from the static environment. RANSAC has no prior knowledge about the historical classification of features; it simply finds those features that fit well to some yet-unknown estimated motion.

The degree of agreement among these measurements regarding the estimated outcome of the PnP algorithm, indicated by the residual reprojection error, is a measure for the amount of trust that can be placed in that particular measurement. Although this residual reprojection error is influenced by various phenomena, such as the quality of the 3D pointcloud constructed by the depth camera, this work is focused on the effects of dynamic environments on this residual reprojection error. More specifically, it analyses to which extent RANSAC effectively protects against these situations by sequentially running the PnP algorithm for different feature sets. A first iteration returns a set of coherent features and a set of outliers. This outlier set might either contain only noise when no dynamic obstacles are present in the scene, or contain another set of coherent features indicating that there was a dynamic obstacle. By running the PnP RANSAC algorithm on this outlier set once more, the motion corresponding to this second coherent measurement set can be estimated and used to generate a second motion hypothesis. To validate this idea, another experiment was conducted with a known ground truth motion.

In Figure 6a, a frame from another real experiment being processed is shown. The paper with markers is moving up with respect to a stationary camera. The RANSAC scheme tracks the dynamic obstacle instead of the static environment. However, within the discarded outliers, another PnP solution can be found based on the features shown in Figure 6b. As to be expected, the calculated motion estimates are quite different for both hypotheses:

$$t_{PNP_1} = \begin{bmatrix} -0.03679 \\ -0.06321 \\ 0.01726 \end{bmatrix}, \quad t_{PNP_2} = \begin{bmatrix} -0.01075 \\ 0.00517 \\ 0.01615 \end{bmatrix} \quad (7)$$

in which t_{PNP_i} is the estimated translation in meters corresponding to the i th generated hypothesis. Hypothesis 2, based on the features in Figure 6b, shows an estimated motion that is closer to the stationary ground truth, especially in the direction of the movement of the paper (y). This is to be expected, as hypothesis 2 is based on actual static features. The next section elaborates on resolving the multiple hypotheses that are potentially generated by the sensor-level IM scheme.




Figure 6. (a) Shows the inliers (green) and outliers (red) for a motion estimated by plain PnP RANSAC as implemented by default, forming hypothesis 1. The outliers of this initial epoch can be used as the input for another iteration, which results in hypothesis 2 with inliers shown in (b).

3.2. SLAM System Integrity Monitoring

The fusion-level IM algorithm detects faults on the sensor fusion scale by investigating to which extent different sensor modalities agree upon the estimated motion of the platform. This is done by leveraging the measurement redundancy in the factor graph, which may also be established by loop closures. Note that, for applications with a critical time-to-alert,

a loop closure might not occur quickly enough, so additional sensor modalities would be necessary. The residuals are defined as the difference between the measured value and the estimated value that a factor produces:

$$q^2 = \sum_i \|h_i(X_i) - z_i\|_{\Sigma_i}^2 \quad (8)$$

The metric q^2 resembles the chi-squared test statistic, commonly employed in integrity monitoring studies to assess the validity of a measurement set. To do this, accurate determination of covariance matrices Σ_i is crucial, as underestimation or overestimation can lead to increased false alarms or missed detections, respectively. Although this paper does not focus on accurate covariance estimation and the scalar value of q^2 , it still indicates whichever of a set of hypothesized trajectories matches the observed measurements best.

4. Simulated Experiments

Through a series of simulated experiments, the sensor-fusion-level integrity monitor is highlighted. A moving platform starts at the origin and traverses the reference trajectory shown in Figure 7a in a counter-clockwise direction. When the platform revisits the starting point, it successfully detects a loop closure.




Figure 7. (a) Shows a simulated reference trajectory of a platform starting at the origin and moving in counter-clockwise direction, following the arrows. (b) Shows the location of a dynamic obstacle fault at the exclamation mark and illustrates the faulty (red) and correct (green) trajectory hypotheses. (c) Shows the corrected trajectories after incorporating the loop closure and how the red trajectory at the location of the dynamic obstacle fault deforms.

4.1. Fault Detection for Pure VO Navigation

In the first experiment, a dynamic obstacle is observed by the moving platform, which would mislead a singular RANSAC filter. The sensor-level IM scheme as discussed in Section 3.1 detects that two hypotheses can be generated based on this dynamic scene. A copy of the original factor graph is generated to keep track of this second hypothesis:

As long as no loop closure is detected, no measurement redundancy is established and the global fusion-level IM can not point out which hypothesis is best, as q^2 is zero for both hypotheses shown in Figure 7b. However, when a loop closure occurs, measurement redundancy is established and the fusion-level integrity monitor finds that hypothesis 1, which includes the faulty visual factors, now returns a non-zero value for weighted residual q^2 , while the graph for hypothesis 2 still has a residual score of zero, as measurement noise is not modelled in this simulated experiment. Recall that the actual value of the score

is not meaningful yet; it is only used to determine which of the two hypotheses is most likely. In this way, this method returns a solution that all observed sensor measurements maximally agree with. Note that this strategy of solving multiple factor graphs in real time is feasible, as optimizing a graph consisting of 360 poses and 361 links takes approximately 133 milliseconds on a typical PC platform.

4.2. Fault Detection for GNSS/VO-Fusion Based Navigation

A second experiment was conducted in a static environment where VO works well but GNSS measurements are affected by a linear building positioning error in the positive y-direction. In this experiment, there are no multiple hypotheses, but the global measurement redundancy caused by sensor multimodality is leveraged. Figure 8a illustrates that, before the loop closure, the optimization residuals are already non-zero, but it is not yet possible to determine which sensor is producing faulty measurements. After the detected loop closure, shown in Figure 8b, the IM may conclude with some certainty that GNSS measurements are off since the loop closure is heavily contradicting the last GNSS pose. Still, a combination of VO drift and a false loop closure might also yield this result. By adding more sensor modalities such as a LiDAR or IMU, fault identification and exclusion can be done with increasing certainty.




Figure 8. Simulated experiment based on the fusion of faulty GNSS (gray dots) and correct VO measurements. (a) Shows the estimated trajectory and the weighted residuals per factor right before the loop closure. (b) Shows the drastically increased value of the q^2 statistic right after the loop closure.

5. Conclusions and Discussion

This work demonstrates a framework for integrity monitoring for graph-based SLAM systems to address critical challenges posed by dynamic environments. Through real and simulated experiments, it is shown how multiple hypotheses can be generated based on the identification of subsets in the complete feature measurement set and how these hypotheses can be resolved, using a fusion-level integrity monitoring algorithm. Since the proposed framework is based on factor graphs, it is sensor-agnostic. This is crucial for high-integrity and accuracy-demanding applications, where a multimodal sensor fusion is required to protect against the surplus of fault modes that can be expected in urban environments with potentially present interference. Certain remaining limitations to be addressed in future work are as follows:

- Accurate estimation of the VO noise model by adequate estimation of the covariance matrix, which is needed to calculate mathematically funded protection levels.
- Estimation of the direction and magnitude of spoofed GNSS signals to correct historical GNSS data for optimal position estimation.

- Incorporation of a feedback loop from the fusion-level IM to the sensor-level IM. When inconsistencies are detected on a global scale, the local sensor-level IM could be tasked to specifically investigate the conflicting measurements in a separate thread.

Author Contributions: Conceptualization, S.B. and H.E.; methodology, S.B. and H.E.; software, S.B.; validation, S.B.; formal analysis, S.B.; investigation, S.B.; resources, S.B.; data curation, S.B.; writing—original draft preparation, S.B.; writing—review and editing, H.E.; visualization, S.B. and H.E.; supervision, H.E.; project administration, H.E.; funding acquisition, not applicable. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets presented in this article are not readily available because the data are part of an ongoing study. Requests to access the datasets should be directed to the corresponding author.

Acknowledgments: Data were recorded using the payload developed by Kjeld van der Linden at Royal NLR.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Joerger, M.; Chan, F.; Pervan, B. Solution Separation Versus Residual-Based RAIM. *Navigation* **2014**, *61*, 273–291. [[CrossRef](#)]
2. Diesel, J.; Luu S. GPS/IRS AIME: Calculation of thresholds and protection radius using chi-square methods. In Proceedings of the 8th International Technical Meeting of the Satellite Division of the Institute of Navigation, Palm Springs, CA, USA, 12–15 September 1995; pp. 1959–1964.
3. Ye, Q.; Gu, Y.; Li, L.; Du, F.; Li, R. Integrity Monitoring for GNSS/INS Integrated Navigation Based on Improved AIME. In Proceedings of the China Satellite Navigation Conference 2024 Proceedings, Jinan, China, 22–24 May 2024.
4. Wen, W.; Pfeifer, T.; Bai, X.; Hsu, L.T. Factor graph optimization for GNSS/INS integration: A comparison with the extended Kalman Filter. *Navigation* **2021**, *68*, 315–331. [[CrossRef](#)]
5. Wen, W.; Meng, Q.; Hsu, L.T. Integrity monitoring for GNSS positioning via factor graph optimization in urban canyons. In Proceedings of the 34th International Technical Meeting of the Satellite Division of the Institute of Navigation (ION GNSS+ 2021), Online, 20–24 September 2024; pp. 1508–1515.
6. Engwerda, H.; Snijders, M.; Casals Sadlier, J. Experimental results of integrity monitoring for UAV flights in urban environment leveraging image based masking. In Proceedings of the 2023 International Technical Meeting of the Institute of Navigation, Long Beach, CA, USA, 24–26 January 2023; pp. 413–427.
7. Snijders, M.; Engwerda, H.; Fidalgo, J.; Domínguez, E.; Moreno, G.; Buendia, F.; Duque, J. P.; Martínez, J.; Martini, I.; Sgammini, M.; et al. Advanced Receiver Autonomous Integrity Monitoring (ARAIM) for Unmanned Aerial Vehicles. *Eng. Proc.* **2023**, *54*, 46. [[CrossRef](#)]
8. Kaess, M.; Ni, K.; Dellaert, F. Flow separation for Fast and Robust Stereo Odometry. In Proceedings of the 2009 IEEE International Conference on Robotics and Automation, Kobe, Japan, 12–17 May 2009; pp. 3539–3544.
9. Xia, X.; Wen, W.; Hsu, L.T. Integrity-constrained Factor Graph Optimization for GNSS Positioning in Urban Canyons. In Proceedings of the 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS) Monterey, CA, USA, 24–27 April 2023; pp. 414–420.
10. Labbé, M.; Michaud, F. RTAB-Map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation. *J. Field Robot.* **2019**, *36*, 416–446. [[CrossRef](#)]
11. Bradski, G. The OpenCV Library. *Dr. Dobbs' J. Softw. Tools* **2000**, *120*, 122–125.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.