FDE Implementations for a Low-Cost GPS/INS Module

Benjamin J Clark and David M Bevly
Auburn University

Biography

Benjamin Clark received M.S. and B.S. degrees in Mechanical Engineering from Auburn University and a B.S. in Mathematics from Freed-Hardeman University. He has worked in the Auburn Mechanical Engineering department's GPS and Vehicle Dynamics Laboratory since 2004 researching GPS/INS integration methods and GPS errors.

David M. Bevly received his B.S. from Texas A&M University in 1995, M.S. from Massachusetts Institute of Technology in 1997, and a Ph.D. from Stanford University in 2001 in Mechanical Engineering. He joined the faculty of the Department of Mechanical Engineering at Auburn University in 2001 as an assistant professor. Dr. Bevly's research interests include control systems, sensor fusion, GPS, state estimation, and parameter identification. His research focuses on vehicle dynamics as well as modeling and control of vehicle systems. Additionally, Dr. Bevly has developed algorithms for navigation and control of off-road vehicles and methods for identifying critical vehicle parameters by using GPS and inertial sensors.

Abstract

This paper describes the implementation details of fault detection and exclusion (FDE) methods for possible use in a low-cost GPS/INS module. The use of a coupled system has been previously considered for improving navigation in GPS difficult environments. Such environments adversely affect the satellite signals received by a navigation system. Integration with an INS allows for higher output rate as well as improved coasting through GPS signal blockage. Previous work has shown that the inclusion of a FDE algorithm allows for the removal of faulty GPS measurements that tend to corrupt the navigation solution. This work gives analysis of several FDE methods both from the standpoint of performance and efficiency.

The resulting goal of this work is the details of a real-time GPS/INS module with FDE improvements. The module design requirements are specified for use in low-cost applications. Since operation in vehicular environments is desired, meter-level accuracy is investigated so that approximate lane-level information could be available to the user of the navigation system. The FDE methods under consideration both have snapshot and sequential implementations. The snapshot methods are performed independently between measurement epochs and thus do not suffer from undetected errors that corrupt the states. Sequential methods are able to detect a wider variety of errors but are delayed in acting on failure conditions. These two implementations are compared for the normalized innovation method and the direct consistency check method. The normalized innovation technique is a comparison of the resulting new information provided by a measurement to a normalized threshold. The threshold is set to detect measurements that do not statistically conform to the expected accuracy. The direct consistency check algorithm performs a comparison of a measurement to what the measurement is expected to be given the removal of the measurement from the estimation. Faults are then detected and removed when inconsistencies are found. These methods are considered for use in the GPS difficult environments.

The comparison of these methods is accomplished by the design and assembly of the GPS/INS module. This module is then used to log the required data for initial post-processing. Various data collection locations considered include open-sky, urban canyon, and heavy foliage areas. The integration and FDE algorithms are then run on the same data sets and fault detection occurrences compared among the methods. Processing time is also monitored for the post-processing to generate efficiency results. The details are then given to implement the chosen method in real-time on a low-cost GPS/INS module.

Due to the quickly changing nature of the GPS errors in difficult environments, the snapshot methods tend to provide faster detection of errors and thus improved performance. For faster implementation, the normalized innovation technique is selected to reduce load on the embedded

navigation system. This choice allows for more flexibility in extending the module use.

The result of this work is a navigation system implementable in real-time that provides improved positioning in GPS difficult environments. Many applications such as vehicle navigation and control benefit from improved performance in these situations. The inclusion of the low-cost requirement allows for more ubiquitous use of these results.

1 Introduction

The difficulty of navigation for a ground vehicle with GPS alone in certain environments is a well known problem. Figure 1 shows an example of navigating with a GPS standalone receiver in a rural neighborhood with moderate foliage coverage. Position jumps of tens to hundreds of meters can be experienced in these situations which makes consistent and continuous positioning impossible. Similar effects occur in urban canyon situations. This research is a continuation of the investigation of improved positioning in these environments.

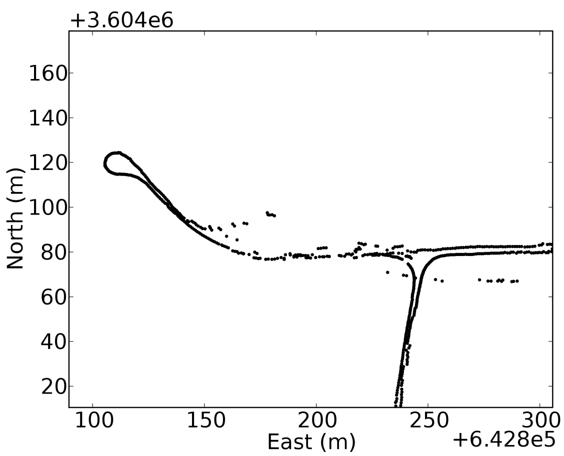

Figure 1. GPS-only Positioning in Foliage

The coupling of an inertial navigation system (INS) with a global positioning system (GPS) receiver inherently gives more robustness to a navigation solution when operating in difficult GPS environments such as heavy foliage and urban canyons. In this paper, these environments will be referred to as shadowed since objects between the receiver antenna and broadcasting GPS satellites cause blockage, reflection, or attenuation of the signals. In these cases severe shadowing frequently make GPS-only navigation impossible. Even when a GPS solution is available, sometimes individual measurements are corrupted due to multipath or attenuation so that the overall navigation solution is degraded. Much of this difficulty is mitigated through the use of an INS which tracks the change in navigation state. Thus, blending together yields a filtered result that reduces the effects of the shadowed environment on the navigation solution.

However, GPS/INS integration can still benefit from a fault detection and exclusion (FDE) method especially in shadowed environments. Since large unmodeled errors can propagate quickly through the solution of the filtered navigation system, the removal of these measurements will provide a more robust solution. There are trade-offs in this case as the removal of measurements may lead to a drifting solution. It has been shown in previous work by the authors that inclusion of a simple FDE routine provides a more robust solution for a closely coupled GPS/INS integration.

Since improved positioning is always desired, this work provides analysis of methods to gain more robust navigation solutions. This is key to the continued and prolific use of GPS in shadowed environments. Since a majority of ground vehicles operate at least part of the time in shadowed environments, improved navigation operation is required. Position-critical applications can also benefit from improved robustness in difficult environments. Also, the use of these algorithms allows for implementation on low-cost modules, such as the target of this research.

1.1 Objectives

The ultimate goal of this research is a low-cost real-time navigation module made of a closely coupled GPS/INS combination made more robust with FDE robustness improvement. Target applications for this module are ground vehicles with lane-level positioning requirements. These vehicles are meant to operate seamlessly in both benign and shadowed environments. To remain low-cost, the prototype module should cost less than $600. Currently the target hardware has been selected to be an embedded package, and thus will have low processing capabilities. The module under development uses the Gumstix microprocessor for integration. This processor allows for a low-cost and low-power module to be built. With the use of a general purpose processor, further additions and improvements can be easily made. The GPS receiver selected for this research is the u-blox LEA-4T. This module provides the raw observables and ephemerides which are necessary to run the FDE methods. A MEMS IMU keeps the overall cost of the system down. A custom board with Analog Devices accelerometers and gyroscopes is currently in development but was not available at the time for this paper. Therefore a Crossbow IMU440 was used which is expected to give similar performance. Through the use of FDE, accuracy bounds are also desired for the module that will provide the user with a measure of trust for the system. The majority of this research focus is on the operation of the system in shadowed environments. As part of this research, performance and processing details are provided for the different FDE algorithms.

1.2 Previous Work

The authors have completed work showing improvemed positioning in foliage environments with the use of a snapshot FDE method [Clark and Bevly, 2008]. In [Oh et al., 2005], a sequential detection method is used to improved the positioning performance of GPS aiding of an IMU. In [Kuusniemi et al., 2002] a snapshot residuals method was used in the positioning of an assisted GPS without an IMU. As was stated in that work, the fault detection methods are more important for positioning with a high sensitivity GPS receiver. This work confirms this statement. A large body of work has been done in the domain of FDE ([Parkinson and Axelrad, 1998], [Sturza, 1998], [Brown, 1998]). Much of this work has been performed for the detection of satellite failures rather than signal failures, which is the focus of this research.

2 Algorithms

Since shadowed environments produce errors that often affect single satellite ranging signals, a closely coupled approach is used rather than a loosely coupled method. Figure 2 describes the flow of information in this method. The navigation processor uses ranges and range rates reported by the GPS receiver to make the position, velocity, and time (PVT) estimation while blending with the INS. This allows the system to perform checks on the individual measurements and thus remove fautly ones. The difference between closely coupled and tightly coupled integrations is that aiding of the ranging processor indicates a tightly coupled framework. Although this nomenclature differs from some current implementations, it is used to distinguish this method from tighter integrations.

2.1 Closely Coupled GPS/INS

The navigation filter is used as a base for the GPS/INS system. The inclusion of an INS provides many advantages for the navigation solution. It allows for a higher update rate driven by data, and thus a theoretically more accurate state propagation between GPS measurements. The states under consideration in the navigation filter are given as

$$\bar{x} = \begin{bmatrix} \psi \\ v \\ x \\ b_a \\ b_g \\ b_c \end{bmatrix} \tag{1}$$

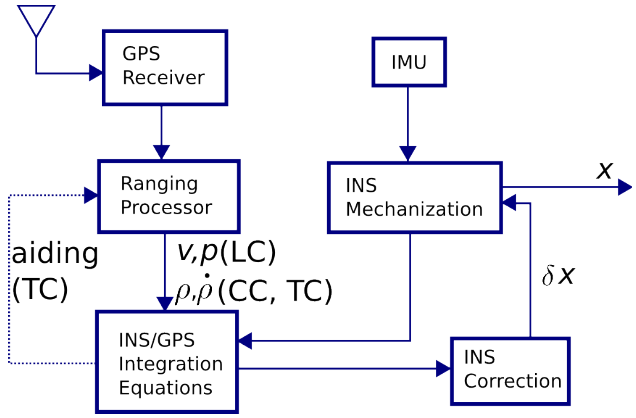

Figure 2. Comparison of Loosely Coupled (LC), Closely Coupled (CC), and Tightly Coupled (TC) Methods

where ψ represents the three-component attitude errors, v are the three components of the user velocity, x are the three components of the user position, b_a, and b_g are the accelerometer and gyro biases, respectively, and b_c is the user clock bias and drift. In this research, the biases were taken to be constant with process noise injected into the time update covariance to keep from holding a previously estimated value and ignoring the data. The attitude, velocity, and position are taken as error states while the biases are used as the true values in this implementation. The state propagation is driven by the INS system's accelerometer and gyro measurements. These equations are described in [Groves, 2008].

When GPS pseudorange and Doppler measurements are available, the system performs a measurement update. For this update, the measurement variances are calculated as functions of the carrier to noise ratio C/N_0. These functions can be found in [Clark and Bevly, 2008]. With these measuremnts, the FDE algorithms are calculated and the valid measurements are applied to the state to continue the operation of the navigation filter.

2.2 Filter Initialization

Methods exist for initializing attitude with a multi-antenna GPS system [Lu, 1995], but there is no intrinsic way to initialize attitude with a single GPS antenna. Alternative sensors such as multi-antenna GPS systems or magnetometers mitigate this difficulty. Initialization is important since the EKF equations are highly nonlinear and dependent on the rotation matrix from the earth-centered earth-fixed (ECEF, e) frame to the body frame (b). Figure 3 illustrates the initialization process.

As data is being processed, the error and bias states are initialized to zeros. When a raw GPS message is available, the pseudoranges, Dopplers, and ephemerides are used

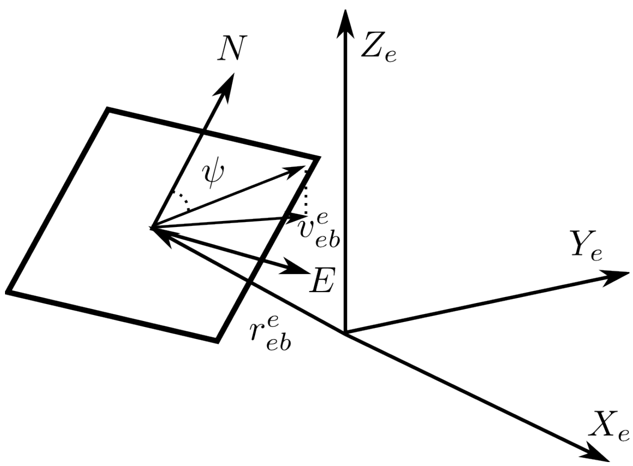

Figure 3. Filter Initialization Diagram

to calculate the current position velocity and time (PVT). With this information, the position and clock terms can be initialized. However, without sufficient velocity, attitude cannot be initialized. Once a velocity threshold has been surpassed, the position solution is used to generate the tangent plane and its rotation matrix from the ECEF plane C_{et}. The velocity is calculated in the ECEF frame so the tangent components (North, East, Down) are calculated using this matrix. With the assumption that the vehicle is operating without a significant vertical component, the heading angle ψ is calculated. This angle is all that is needed to fully define the ECEF to body frames as

$$C_{eb} = C_{et} \begin{bmatrix} \cos\psi & -\sin\psi & 0 \\ \sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2}$$

With this information, all the states are initialized and data processing can proceed.

2.3 Snapshot Normalized Innovation

With the state covariance representing the expected variances on the states and the measurement covariance representing the expected variances of the measurements due to signal attenuation and noise (which are functions of the C/N_0 and receiver parameters), a normalized innovation can be calculated as a test statistic for each measurement to determine whether or not the new information fits into the expected range of values. The calculation of each normalized innovation takes the form of

$$y_i = \frac{z_i}{\sqrt{C_{ii}}} \tag{3}$$

Where $C = HPH^T + R$, which is calculated as part of the filter equations.

These normalized innovations should (under normal operating conditions) be normally distributed with zero mean and unit variance. Assuming the normalization makes the statistic unit variance, values that lie outside the threshold indicate that the innovations are non-zero mean and thus have errors that would bias the navigation solution. This case is shown in Figure 4. These measurements are to be exclueded from the measurement update.

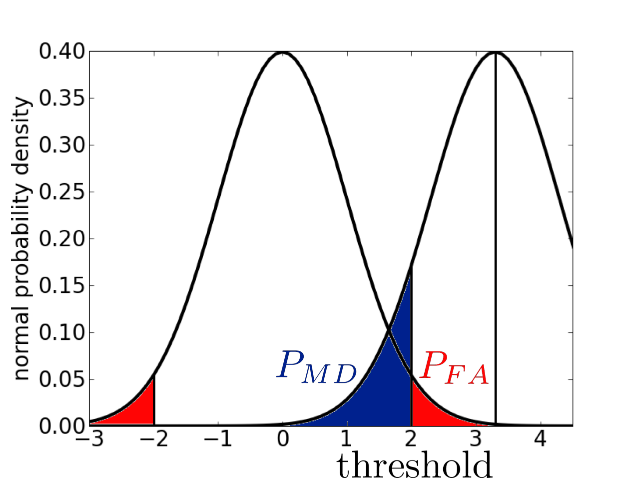

Figure 4. Normalized Innovation Density Function for Valid and Invalid Measurements

Selection of the threshold amounts to determining the probability of false alarm P_{FA} that is tolerable for the application. With a given P_{FA}, the threshold is selected so that the area under a zero mean unit variance Gaussian distribution from the threshold toward positive and negative infinity equals the P_{FA}.

In the event that a pseudorange is rejected, the corresponding pseudorange rate measurement should also be rejected. However, a rejected rate measurment does not require the removal of the corresponding pseudorange [Groves, 2008].

2.4 Sequential Normalized Innovation

The sequential monitoring of the normalized innovation is also called innovation sequence monitoring [Groves, 2008]. These methods allow for the detection of slower building discrepancies in the innovations that a snapshot method could not detect. Taking the last N measurements into account, the average performance of an innovation can be estimated. The test statistic for this method is

$$\mu_{kj} = \frac{1}{N} \sum_{i=k+1-N}^{k} y_{i,j} \tag{4}$$

Taking the mean of N zero-mean unit variance Gaussian distributed random variables gives standard deviation of $\frac{1}{\sqrt{N}}$. This implies that the test statistic is distributed with zero mean and $\frac{1}{N}$ variance. Therefore the threshold can be scaled by $\frac{1}{\sqrt{N}}$ for detecting with a consistent P_{FA}.

2.5 Snapshot Direct Consistency

The direct consistency check methods have been described in the literature [Parkinson and Axelrad, 1998]. Alternative ways to accomplish this check have been shown to be equivalent in their ability to detect faults [Brown, 1998]. The main idea behind these methods is to apply the received measurements in the measurement update and then check the reiduals after this update. An advantage of this method is that it can be performed in a least-squares fashion to get an independent (non-filtered) solution of position and velocity with enough satellites in view. The main difference between this method and the normalized innovation method is that here the FDE occurs after the measurement update is applied. This method ignores the system's ability to propagate the states and only checks the nature of the post-measurement update residuals.

Since the residuals are not independent after the update, the test statistic

$$s^2_{\delta z,k} = \delta z^T C^{-1} \delta z \tag{5}$$

has a chi-square distribution with fewer degrees of freedom than there are measurements. This requires satellite coverage sufficient to generate a distribution with multiple degrees of freedom [Brown and Chin, 1998].

The chi-square distribution changes as the number of satellites increases through the distribution's degrees of freedom. Again since an erroneous measurement has an unknown bias, the test amounts to comparing a centralized and noncentralized chi-square distribution and selecting the amount of P_{FA} that is tolerable. Examples of these distributions are shown in Figure 5.

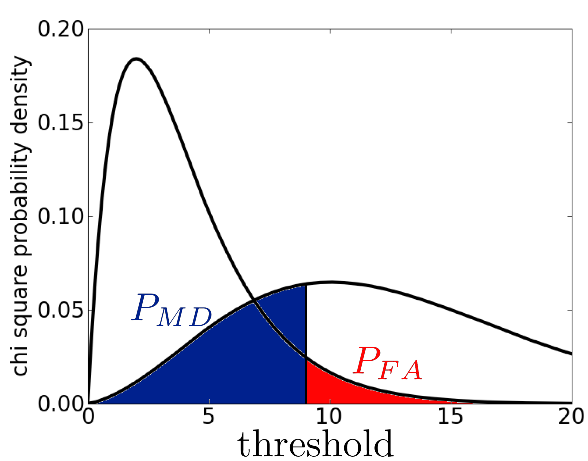

Figure 5. Direct Consistency Check Distibution for Valid and Invalid Measurements

2.6 Sequential Direct Consistency

The direct consistency check method can be extended by using a test statistic that sums the previous residuals for a pre-determined window size N. Thus, the test statistic becomes

$$l = \sum_{i=k-N+1}^{k} \delta z_i^T C_i^{-1} \delta z_i \tag{6}$$

This statistic provides improved detection capabilities at the expense of delayed alert time.

Summing chi-square distributions yields a chi-square distribution with a different number of degrees of freedom. The resulting distribution degrees of freedom equal the sum of the individual degrees of freedom. Therefore, the selection of the threshold is equivalent to that of the snapshot direct consistency check but with a different number of degrees of freedom for the distribution.

3 Methodology

To compare the methods under consideration, the same data is run through all four methods for all data under consideration. Each method has a threshold which is a tradeoff between the method's ability to detect errors and the false alarm rate tolerable for the application. The implemented detection algorithms will generate test statistics for each of the areas where data is logged. The thresholds for these test statistics were selected for a consistent probability of false alarm, P_{FA} of 0.3%. This was the value used in [Clark and Bevly, 2008] and used to compare performance to previous work. The speed of the detection algorithms is also considered due to the increased computational complexity of some of the algorithms. In this case, increased storage needs for sequential methods is neglected except for the necessity in accessing multiple memory locations.

3.1 Hardware Module

Logging of the required raw data came from a u-blox LEA-4T GPS receiver and a Crossbow IMU440 inertial system. The logging was performed by a Gumstix Verdex Pro computer. This computer is the target for the real-time implementation to be the continuation of this work. These tests were performed in an Infiniti G35 used by Auburn University's GPS and Vehicle Dynamics Laboratory. Processing of the raw data was accomplished on a PC for simultaneous comparison of the fault detection methods.

3.2 Logging Location

The path selected for logging was chosen to include both open sky situations and heavy foliage for comparing the test statistics. Figure 6 shows a sky view of the area logged with the number of satellites visible along the path indicated by color. As color goes from red to violet, the number of satellites goes from 2 to 10. Certain sections of this path are shown to be impossible to operate with a GPS-only system due to limited satellite visibility.

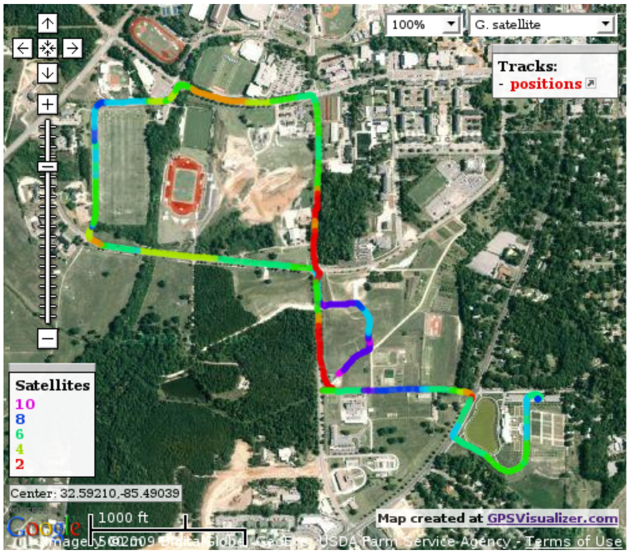

Figure 6. Overhead View of Logging Path

4 Results

As was described earlier, the fault detection and exclusion methods are more beneficial for higher sensitivity recievers. Even in the benign environment, fault detection is difficult with lower sensitivity receivers since gross signal errors from the local environment cause loss of lock rather than degraded measurements. Therefore the reciever will not be able to report measurements that could be detected. Instead, they are not available at all since tracking fails for them. This was the case for the GPS/INS module under consideration. Although the path shown in Figure 6 has areas of moderate tree coverage, tracking was expected to be possible in these areas.

Figures 7 and 8 show the detection performance of the four methods each scaled by their threshold for comparison. The scaling allows for easier comparison of the methods. Thus, a fault is detected when any of the test statistics cross the threshold with magnitude of 1.

As was described, each satellite has separate range and range rate measurements and therefore separate normalized innovation test statistics. The residual test takes the system operation as a whole and attempts to detect faults by combining all residuals into a single test statistic.

Table 4 shows the average time to perform each of the FDE methods. The snapshot methods take less time than their sequential alternatives and residual methods take more time in general than normalized innovation due to

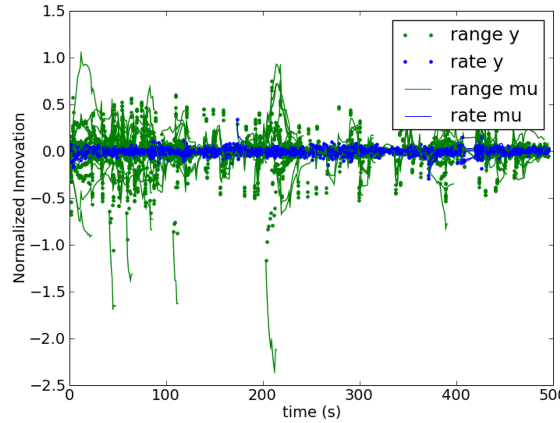

Figure 7. Normalized Innovation Test Statistics for Path

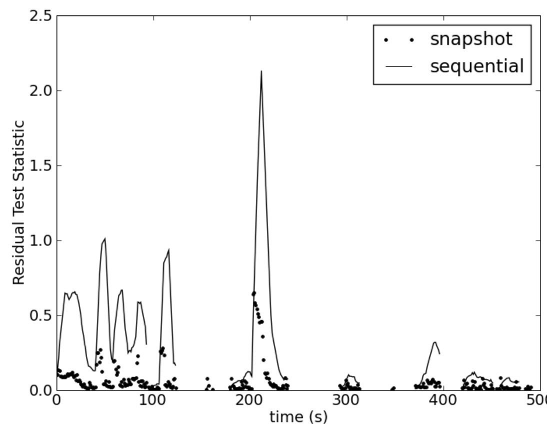

Figure 8. Residual Test Statistics for Path

the necessity of performing the measurement update before the test statistic calculation. Although performance is not much different, some gain is seen with the snapshot normalized innovation method.

FDE Method	Average time (ms)
Snapshot Normalized Innovation	0.149
Sequential Normalized Innovation	0.472
Snapshot Direct Consistency	0.382
Sequential Direct Consistency	0.798

An example fault detection occurred at 45 seconds into the run, as shown in Figures 9 thru 11. The path in Figure 9 is shown with both monitored and corrected paths shown. In the monitored case, the test statistics were calculated but no fault exclusion took place. The corrected path shows the effect of rejecting the detected faulty measurement. This point in the path is along a hill with a few tall trees blocking low elevation satellites. Since the cover is sparser here, the signal was not lost as at other times but instead included an erroneous component. The normalized innovation methods detected the fault in this new signal but the residual method never got above the threshold to indicate detection.

Figure 9. Position Improvement Using FDE Method

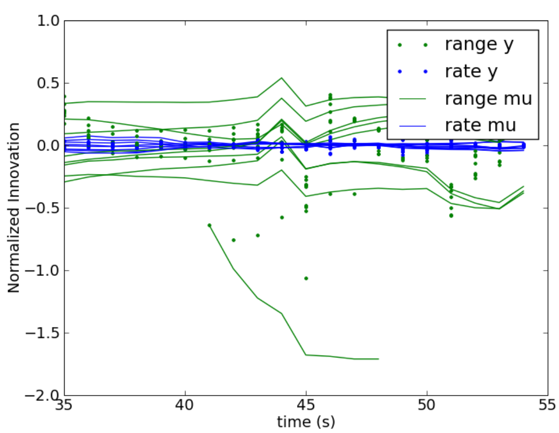

Figure 10. Normalized Innovation Test Statistics for Corrected Path

5 Conclusions

The comparisons given in this paper showed the use of multiple statistical detection methods for the same data set. The performance of these methods was also compared for implementation in real time on a low-cost navigation unit. The resulting analysis shows the benefit of the snapshot normalized innovation method due to good detecton performance and higher efficiency in this implementation. The sequential methods did not allow for sufficient time to detect errors in all cases and thus could not be justified with the increased processing time necessary for their use. The resulting system is capable of running real-time with good performance and has great potential in the area of low-cost navigation systems.

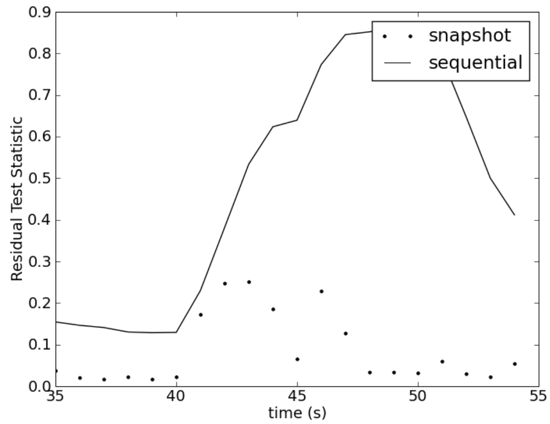

Figure 11. Residual Test Statistics for Corrected Path

References

Brown, R. G. (1998). A baseline gps raim scheme and a note on the equivalence of three raim methods. In *Global Positioning System, Volume V: Papers Published in NAVIGATION*. Institute of Navigation.

Brown, R. G. and Chin, G. Y. (1998). Gps raim: Calculation of threshold and protection radius using chi-square methods - a geometric approach. In *Global Positioning System, Volume V: Papers Published in NAVIGATION*. Institute of Navigation.

Clark, B. and Bevly, D. (2008). Gps/ins integration with fault detection and exclusion in shadowed environments. In *Proceedings of the IEEE/ION Position, Location, and Navigation Symposium*.

Groves, P. D. (2008). *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Artech House.

Kuusniemi, H., Leppakoski, H., Syrjarinne, J., and Takala, J. (2002). Fault detection and isolation in high-sensitivity assisted gps. In *Proceedings of teh ION GPS*.

Lu, G. (1995). *Development of a GPS Multi-Antenna System for Attitude Determination*. PhD thesis, University of Calgary.

Oh, S., Hwang, D.-H., Park, C., Lee, S., and Kim, S. (2005). Attitude determination gps/ins integrated navigation system with fdi algorithm for a uav. *Journal of Mechanical Science and Technology*, 19:1529–1543.

Parkinson, B. and Axelrad, P. (1998). Autonomous gps integrity monitoring using the pseudorange residual. In *Global Positioning System, Volume V: Papers Published in NAVIGATION*. Institute of Navigation.

Sturza, M. (1998). Navigation system integrity monitoring using redundant measurements. In *Global Positioning System, Volume V: Papers Published in NAVIGATION*. Institute of Navigation.