Optimising the algorithm design for high-integrity relative navigation using carrier-phase relative GPS integrated with INS

Paul D Groves, Charles R Offer, Christopher J Mather, Graham W Pulford, Immanuel A Ashokaraj, Alex A Macaulay
QinetiQ ltd, Farnborough, United Kingdom

BIOGRAPHY

All authors are members of the Navigation and Positioning Algorithms team within QinetiQ's UAVs and Autonomous Systems Line of Business.

Dr Paul Groves specialises in integration and modelling of navigation systems. He has over 10 years experience in navigation systems research and development for air, land and marine applications, including with GNSS, inertial navigation, terrain-referenced navigation and various dead-reckoning sensors. He is an author of about 30 technical publications, including the new book, *Principles of GNSS, Inertial and Multi-Sensor Integrated Navigation Systems* (Artech House). He holds a BA/MA and a DPhil in physics from the University of Oxford. He is a fellow of the Royal Institute of Navigation, serving on the R&D and Technical committees, and an active member of the ION.
pdgroves@QinetiQ.com

Dr Charles Offer has been a scientist with QinetiQ since 2004. He has worked on methods of high-integrity relative navigation using INS and carrier phase GNSS, for example for shipboard relative navigation, and on INS/GNSS methods of motion compensation for synthetic aperture radar. He holds an MA MSci and a PhD in physics from the University of Cambridge.

Christopher Mather is a senior scientist within QinetiQ. He joined in 2000 and his primary focus has been on integrated INS/GPS navigation with low cost sensors. He has also been looking at the integrity of GPS-based landing systems for maritime aircraft. He holds a BSc (Hons) degree in mathematics from Southampton University and an MSc in mathematical modelling and numerical analysis from the University of Oxford.

Dr Graham Pulford received the BE and BSc (Hons) from UNSW in 1988 and PhD from ANU in 1992. His research experience includes 7 years with the University of Melbourne Sensor and Signal Processing lab (CSSIP) and 8 years in Industry. His primary research interests are in tracking and data fusion and their application to navigation systems. He has worked on tracking systems for the Australian Jindalee Operational Radar Network and the combat system for the Collins class submarine. He has more than 60 publications including 15 journal papers and a US patent. He is a senior member of the IEEE.

Dr Immanuel Ashokaraj is a navigation systems (algorithm) engineer and has been with QinetiQ since 2007. He received his PhD from DCMT, Cranfield University in 2005 were he worked on low cost navigation for single & multiple vehicles and high accuracy integrated INS/GPS navigation for marine surveying applications.

Alex Macaulay is the Capability Group Leader for the Navigation and Positioning Algorithms Team within QinetiQ. He joined in 1999 and has been the technical lead on a range of navigation programs including the development of a low-cost integrated INS/GPS navigation system for helicopter low-visibility operation and the assessment of integrated INS/GPS navigation techniques for complex weapons. His recent focus has been on GPS-based landing systems concepts for maritime aircraft. He holds a BEng (Hons) in Aeronautical Engineering and a Masters in Information Technology, both from the University of Glasgow.

ABSTRACT

An increasing number of applications require the provision of a relative navigation solution which exhibits both high accuracy and high integrity. Examples include air-to-air refuelling (AAR) of unmanned air vehicles (UAVs) and automated landing of aircraft on ships. In a previous paper, QinetiQ demonstrated that adding inertial navigation

to a carrier-phase relative GPS system can improve performance in a number of respects.

This paper proposes an optimised processing architecture for high-integrity carrier-phase relative INS/GPS. This comprises at least one “single node” navigation processor aboard each vehicle and a relative navigation processor aboard the roving vehicle. For protection against reference user equipment failure, a fusion algorithm may be used to combine information from multiple single-node navigation processors aboard the reference vehicle. A patent-pending technique has been developed to ensure that the double-differenced carrier phase ambiguities remain integer multiples of the wavelength.

The options for how and whether to combine GPS measurements, what to estimate as Kalman filter states and how many satellites the integration algorithms should handle are discussed and preferences identified. The trade-off between a partitioned and a differenced architecture for the relative navigation algorithm is discussed.

The use of parallel solution hypotheses for fault isolation and exclusion is described and the integrity monitoring architecture summarised. Lastly, further work is proposed to improve the integrity of ambiguity fixing and develop a hybrid analytical and simulation-based performance model for determining a lower bound to relative INS/GPS solution availability.

1. INTRODUCTION

One of the most challenging problems in navigation is to provide a relative position solution that is both accurate and has high integrity. Applications that require this include automated air-to-air refuelling (AAR) and automatic landing on ships, of both unmanned air vehicles (UAVs) and manned air vehicles, formation flying, separation assurance for civil aircraft, both in the air and on the ground, and train collision avoidance.

Many future UAVs, particularly those with an offensive capability, will need to operate with mission times of several tens of hours. A key enabler for this level of endurance is the ability to refuel whilst airborne. An important aspect of the automated AAR problem is the ability to deduce position relative to the tanker, a task traditionally achieved by the pilot processing information from the visual scene. A number of design concepts under consideration [1] have sought to replace or augment this visual processing task with a relative navigation capability based on high-integrity relative Global Positioning System (GPS) and inertial navigation. Such a high-integrity relative navigation capability may be used to facilitate the UAV rendezvous and station-keeping with the tanker. It can maintain safe separation between the aircraft whilst additional short-range sensors are used to accomplish the final hook-up with the tanker.

Landing of aircraft on ships requires manoeuvring within a tight space while the ship is continually pitching and yawing with respect to the sea. An erroneous navigation solution could result in a collision, damaging aircraft and ship, endangering crew and preventing other aircraft from landing. Therefore, there is a high integrity requirement – the navigation system must be able to verify that it is fault-free with a high degree of confidence; when it cannot do this, an alert must be raised.

QinetiQ has extensive experience of developing robust integrated navigation solutions to challenging problems. Throughout the 1990s, research was conducted into improving the operational effectiveness of covert automatic air-to-air refuelling aimed at future offensive aircraft. This included integrated flight management system concepts, omni-directional approach and integrity-monitored relative GPS navigation, culminating in a full flight demonstration system. This was closely linked to research into the recovery of large helicopters on to small ships.

In 2005, QinetiQ demonstrated the world’s first automated landing of a short take-off and vertical landing (STOVL) aircraft on an aircraft carrier [2, 3]. In 2007, it contributed to the Broad Agency Announcement (BAA) phase of the US Joint Precision Approach and Landing System (JPALS) program.

Both AAR and automated landing require a sub-metre relative navigation solution between the two vehicles to be guaranteed. In practice, to achieve sub-metre alert limits for the integrity monitoring system, the fault-free accuracy of the relative navigation solution needs to be sub-decimetre. This requires a carrier-phase relative GPS solution.

A further requirement for these applications is continuity. It is not acceptable for the relative navigation solution to fail or an integrity alert to be raised without prior warning as the docking or landing manoeuvre may have reached a critical phase where it cannot be aborted. Therefore, a high probability of the navigation solution remaining usable for a defined period into the future must be maintained. In practice, this means that the system should continue to provide a usable relative navigation solution for a defined time in the event of a fault occurring. Furthermore, it should be able to predict in advance when there is insufficient information for the integrity monitoring to operate.

Continuity is difficult to provide using GPS alone. There are not always sufficient signals or adequate signal geometry to provide an integrity-monitored navigation solution in the case of a signal being rejected due to a fault, though these scenarios can usually be predicted in advance, enabling a continuity alert to be raised. However, it is also possible for the complete GPS navigation solution to suddenly fail due to deliberate jamming, incidental interference or user equipment failure.

Adding an inertial navigation system (INS) to the relative navigation system massively improves continuity performance by ensuring that the navigation solution accuracy degrades gradually in the event of a GPS failure. This enables the docking or landing manoeuvre to be safely completed or aborted. It can also improve the sensitivity of some of the integrity monitors and aid detection and repair of cycle slips, avoiding the need to re-fix the integer ambiguities after brief interruptions of carrier-phase tracking [4].

QinetiQ has extensive experience of advanced INS/GNSS and multi-sensor integrated navigation. Recent work has included the demonstration of non-coherent deep INS/GPS integration for optimised signal-to-noise performance [5] and the development of multi-sensor integrated navigation systems for pedestrians [6, 7] and for underwater vehicles. Earlier work included development of a position and attitude determination system for day/night all-weather helicopter operations, an INS/GPS tracking system for the Williams Formula 1 racing car team, QinetiQ's patented adaptive tightly-coupled (ATC) INS/GNSS integration technique [8], advanced terrain referenced navigation [9] and robust transfer alignment [10].

Individually, meeting the navigation requirements of applications such as automated landing on ships and air-to-air refuelling of UAVs requires three navigation technologies to be combined:

- INS/ global navigation satellite system (GNSS) integration;
- Relative carrier-phase GNSS;
- High integrity GNSS.

Individually, each of these three technologies is well known and the basics are covered in a number of standard text books [11, 12, 13]. However, combining them raises a number of new challenges, including:

- How to combine INS/GNSS integration, which uses a Kalman filter, with carrier-phase ambiguity resolution, which may use a Kalman filter.
- How many of the standard integrity monitoring methods, such as used for the Local Area Augmentation System (LAAS) [14], which assume a known reference station position, may be adapted to relative navigation.
- How the INS may be used to aid some of the GNSS integrity monitoring.
- How to isolate faulty measurement streams from the navigation solution when Kalman filters are used for INS/GNSS integration and GNSS ambiguity resolution, resulting in old measurement data impacting the navigation solution.
- How to model the performance of the position solution when the analytical models applied to a snapshot GNSS solution are not readily extendable to a filtered navigation solution, while using Monte-Carlo simulation alone is computationally unfeasible for high integrity applications. This issue is largely untouched in the literature.

This paper examines the design trade-offs that must be made to optimise the algorithm design of a high integrity carrier-phase relative INS/GPS system. The main issues to consider are processor load, data-link bandwidth, solution accuracy and fault tolerance/ robustness. It also discusses approaches to integrity monitoring and (briefly) navigation system performance modelling.

Section 2 describes QinetiQ's top-level processing architecture. Sections 3 to 5 then discuss the algorithm design trade-offs for the three types of navigation processor: single-node, reference data fusion and relative navigation. Section 6 then discusses the integrity including monitor types and fault exclusion. Sections 7 and 8 discuss further work and present conclusions, respectively.

Note that the term "rover" is used to describe the vehicle that requires the relative navigation solution and "reference" the vehicle that the rover is navigating with respect to. Thus for AAR, the refuelling UAV is the rover and the tanker aircraft is the reference, whereas for shipboard relative navigation, the aircraft is the rover and the ship is the reference.

2. TOP-LEVEL PROCESSING ARCHITECTURE

Figure 1 shows QinetiQ's preferred top-level architecture for high-integrity relative INS/GPS navigation. A centralised implementation of the navigation processing is not considered practical due to the data-link capacity and robustness required to transmit a full set of IMU measurements from the reference vehicle to the roving vehicle. Also, the processor load aboard the rover could be excessive. Instead, it is proposed to distribute the processing between the rover and reference in a type of federated integration architecture [11, 15].

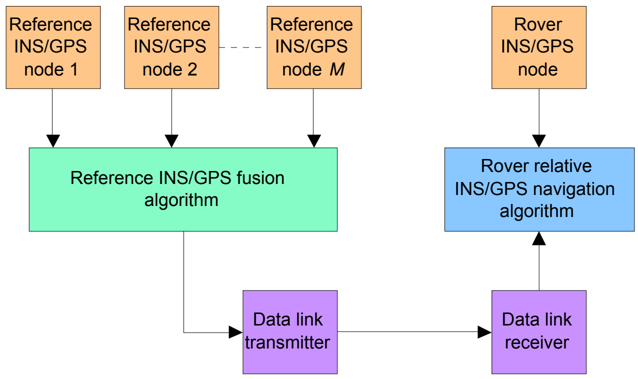

Figure 1: Top-level architecture for high-integrity relative INS/GPS navigation

Each set of GPS user-equipment is accompanied by an inertial measurement unit (IMU), located in the vicinity of the GPS antenna. These output to a single-node navigation processor which performs:

- Inertial navigation computation [11, 16];
- INS/GPS integration [11];
- Float GPS carrier-phase ambiguity estimation;
- Integrity monitoring (see Section 6).

The combination of GPS user equipment, IMU and single-node navigation processor is referred to here as a node. Optimising the single-node algorithm design is discussed in Section 3.

Multiple nodes are proposed for the reference system to protect against user equipment failure. LAAS uses this approach [14]. Note that at least three nodes are needed to be able to detect all user-equipment failure modes (common-mode software failures excepted), identify the faulty node and verify that the remaining nodes are fault-free, a concept known as fault detection and exclusion (FDE) (using the definition in [11]). The reference INS/GPS fusion algorithm then combines data from the reference nodes to produce a virtual INS/GPS navigation solution located at the centroid of the nodes. Although not shown in Figure 1, the same approach could be adopted for the rover. Optimising the reference data fusion algorithm design is discussed in Section 4.

Note that the option of combining a single IMU for the reference vehicle with multiple sets of GPS user equipment has been rejected on two grounds. Firstly, providing only a single IMU provides the reference vehicle navigation system with no robustness against inertial sensor failure. Secondly, a shared IMU would be situated much further away from each GPS antenna than dedicated IMUs would be, resulting in a much larger lever arm between the IMU and antenna. As lever arms are subject to angular flexure and vibration as the host vehicle manoeuvres, larger lever arms lead to larger IMU-antenna lever arm uncertainties. Consequently, inertially-aided GPS integrity monitors would be less sensitive and subject to greater manoeuvre-induced biases.

The relative navigation algorithm determines the relative position and velocity of the rover with respect to the reference using information from both the rover and reference navigation systems. Correlations between the bias-like GPS errors affecting the rover and reference measurements enable a relative navigation solution to be obtained which is more accurate than the rover and reference absolute navigation solutions. In particular, the double differencing of GPS carrier-phase measurements between satellites and receivers enables an integer constraint to be applied to the ambiguities in those measurements. This allows ambiguity fixing to be used to improve precision further. Optimising the relative navigation algorithm design is discussed in Section 5.

The data link transmits fused information from the reference navigation system to the rover by radio for use in the relative navigation processor. This includes:

- Reference system navigation solution;
- Reference system GPS measurement data, including measurement noise model coefficients (for use in the relative INS/GPS Kalman filter), issue of data ephemeris (IODE) and issue of data clock (IODC);
- Tuning information for the relative navigation algorithm;
- Initialisation information for the relative navigation algorithm (including float ambiguities and Kalman filter error covariance information);
- Reference system INS corrections (see Section 5);
- Information for the integrity monitoring system (see Section 6).

The data link must be designed for the optimum trade-off between data rates, data latency and precision, given the capacity available.

Corresponding information is supplied to the relative navigation processor by the rover single-node navigation processor. This does not require a radio link, so capacity constraints and data latency are not a major problem.

3. SINGLE-NODE INS/GPS

Figure 2 depicts a simplified flowchart for each single-node INS/GPS navigation processor, excluding the integrity monitoring. The INS/GPS integration architecture is closed-loop tightly-coupled as defined in [11]. An integration algorithm based on the extended Kalman filter (EKF) inputs GPS pseudo-range and accumulated delta range (ADR) measurements and feeds back position, velocity and attitude corrections, together with IMU error estimates, to the inertial navigation equations processor.

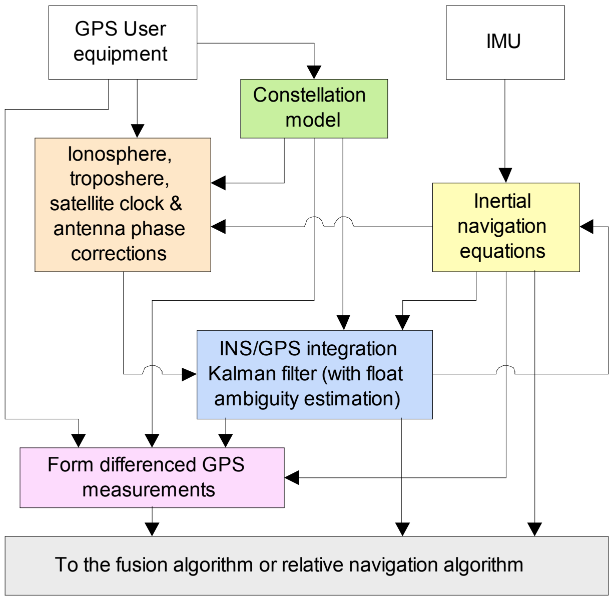

Figure 2: Simplified flowchart for navigation processor in reference and rover INS/GPS nodes (excludes integrity monitoring)

As well as the usual INS and GPS error states (see below), the integration algorithm also estimates (undifferenced) float carrier-phase ambiguities. These ambiguity estimates and their covariance are used to initialise the ambiguity states in the relative navigation filter (see Section 5). This reduces the need for a significant convergence period between initialisation of the relative navigation processor and there being sufficient filter convergence for the ambiguities to be fixed. This is critical in applications where there is only a short period between establishment of the reference – rover data link and the need for a full-precision relative navigation solution. This technique is sometimes known as "pre-filtering".

The main design trade-offs are in how to input the GPS measurements into the integration Kalman filter and which quantities to estimate as states.

Firstly, given that the single-node navigation processors are to estimate float ambiguity states for pre-filtering, code pseudo-range and ADR measurements must be input separately. Otherwise either carrier-smoothed pseudo-ranges or pseudo-ranges, together with pseudo-range rates, Doppler shifts or delta ranges, could be used.

Considering next whether to combine GPS measurements on the L1 and L2 frequencies, assuming Precise Positioning Service (PPS) user equipment, there are two main options for the pseudo-range measurements: separate L1 and L2 measurements or the ionosphere-free combination, which may be smoothed to reduce noise as discussed in [11]. For the ADR measurements, there are three main options to consider: separate L1 and L2 measurements, the ionosphere-free combination and the wide-lane combination (maximising the wavelength to reduce the search space for ambiguity resolution).

The preferred option is to input separate L1 and L2 measurements into the integration algorithm for the following reasons:

- It is robust against signal interruption on one or other frequency;
- It can optimally weight the measurements on the two frequencies according to signal to noise level;
- There is more flexibility in calibrating the ionosphere propagation errors if they are estimated as Kalman filter states;
- Ionosphere calibration information can be shared between code and carrier without having to combine the code and carrier measurements;
- Single-frequency ADR measurements are less noisy than the ionosphere-free and wide-lane combinations.

The main drawbacks are:

- The number of measurements is doubled and the number of states increased, so an EKF requires up to 8 times as much processing power [11, 17];
- Dual-frequency ambiguity resolution also requires more processing power.

If there is insufficient processing power available to implement the preferred option, an alternative to consider is processing ionosphere-free combinations of both pseudo-range and ADR. However, the wavelength of the ionosphere-free combination of carrier measurements is 107 mm, requiring about twice the precision for ambiguity resolution compared to GPS L1 measurements, while the tracking noise is multiplied by a factor of ~3.4 [11].

Pseudo-range and ADR measurements are differenced twice (as distinct from doubly-differenced) prior to their output to the reference fusion algorithm or relative navigation filter. Firstly, they are differenced with ranges estimated from the satellite ephemeris data and single-node navigation solution to eliminate most of the dynamics. Secondly, they are differenced with average values across the measurement set to eliminate the receiver clock errors. Ambiguity shifts are also applied. The net result is to vastly reduce the dynamic range of the GPS measurements so that less data-link capacity is required. Corrections from ionosphere and troposphere models are not applied so that the reference and rover navigation system designers have the option to use different models from each other.

Moving on to state estimation, the preferred option for the navigation and IMU error states is to estimate position, velocity and attitude errors and accelerometer and gyro biases as Kalman filter states on the basis that this is the most common state selection for INS/GPS integration [11]. Aviation-grade (i.e. Standard Navigation Unit (SNU) 84 specification) inertial sensors should be used to reap the full continuity and integrity monitoring benefit from incorporating inertial navigation. Therefore, it may prove viable to omit either the gyro biases or both the accelerometer and gyro biases in order to reduce the processor load, while estimating higher-order IMU errors is unlikely to bring significant benefit.

For the GPS states, ionosphere propagation delays must be estimated if L1 and L2 measurements are processed separately by the EKF, while for pre-filtering, one ambiguity state per ADR measurement must be estimated. The other states to consider estimating are the receiver clock bias and drift and the range biases.

The principal reason for estimating the receiver clock bias and drift is to enable the measurement timing data provided by the GPS user equipment to be corrected so that data from the various reference navigation system nodes and the rover navigation system may be properly time-synchronised. This is important if the GPS user equipment does not correct its timing outputs using its internal navigation processor.

If receiver clock states are not estimated all GPS measurements and Kalman filter states must be differenced between satellites; this reduces the processor load.

Range biases model the residual troposphere and satellite clock errors, after the application of corrections, together with the residual line of sight ephemeris errors for each satellite tracked. It is preferable to estimate them as Kalman filter states for integrity reasons (see Section 6). Including range biases gives a more representative error covariance matrix, aiding determination of protection levels by the integrity monitoring function. They also prevent the range biases from biasing the measurement innovations, enabling more effective innovation-based integrity monitoring. The main drawback of range bias estimation is increased processor load due to the increased number of Kalman filter states. Note also that range biases are only partially observable as there is insufficient measurement information to fully separate them.

The final single-node design issue to address is how many satellites the integration algorithm should handle. Best performance is obtained by supporting 12 satellites, the maximum number that most GPS user equipment can track. However, reducing the number of satellites handled to between 8 and 10 should only have a small impact on performance, but will reduce the processor load by 30-70%. Where only some GPS measurements are processed, those from the highest elevation satellites should be selected as these are least susceptible to multipath and signal blockages and they exhibit lower ionosphere and troposphere propagation delays. Reducing the number of satellites handled to 6 or fewer should not be considered as this is expected to significantly reduce the ability of the navigation processors to use consistency information for ambiguity resolution and detection of faulty satellite signals.

4. REFERENCE DATA FUSION

Fusing the data from the reference navigation system nodes prior to transmission over to the data link has a number of benefits. Firstly, less information needs to be transmitted for the main relative navigation solution, enabling data transmission lags to be reduced. Note that the additional information is still needed to maintain integrity (see Section 6), but can be transmitted with a greater lag.

A significant benefit of working with a fused-reference-data navigation solution in the rover is that fewer carrier-phase integer ambiguities need to be resolved. Ambiguity fixing is one of the most processor-intensive aspects of carrier-phase relative GPS. Therefore running a single ambiguity fixing algorithm with fused reference data brings significant processing efficiency benefits over running separate ambiguity fixes with data from different reference system nodes.

A further benefit of fusing the reference navigation solution is that it removes the need for the relative navigation processor within the rover to model the effects of flexing lever arms between the reference system nodes.

However, taking a simple average of ADR measurements from multiple INS/GPS nodes changes the ambiguities from integer multiples of the carrier wavelength, λ_{ca} (when double-differenced), to integer multiples of λ_{ca}/M, where M is the number of nodes fused. Assuming Gaussian noise uncorrelated between nodes, averaging improves the precision of the float ambiguities by a factor of $M^{1/2}$. If there are N ambiguities to be fixed within the relative navigation algorithm, the number of candidate ambiguity vectors to search is increased by a net factor of $M^{N/2}$. Thus, averaging of ADR measurements requires either vastly more processing power to be allocated to the ambiguity fixing aboard the rover vehicle and/or more data to be gathered by the float ambiguity filters (embedded here in the INS/GPS integration filters) prior to fixing. Neither of these options is necessarily practical.

QinetiQ's patent-pending [18] solution to this problem is to fix the relative ambiguities between the reference system nodes within the reference data fusion algorithm. If the inter-node relative ambiguities are known, corrections can be applied to the averaged ADR measurements transmitted over the data link such that the ambiguities remain as integer multiples of λ_{ca}.

One form of ambiguity correction is

$$\delta a_{n,r}^{c} = \frac{1}{M}\sum_{i=2}^{M} \nabla\Delta\breve{a}_{n,r}^{i,1}, \quad (1)$$

where $\nabla\Delta\breve{a}_{n,r}^{i,1}$ is the fixed ambiguity double-differenced between reference nodes i and 1 and between satellites n and r, where r denotes the common reference, and the superscript c denotes the reference centroid. This is equivalent to adjusting the ADR measurements from nodes 2 to M so that their ambiguities are the same as those of node 1.

Figure 3 depicts a simplified flowchart for the reference data fusion processor. The time synchronisation and sub-sampling function synchronises the times of validity of the GPS measurements from the different reference nodes and then synchronises all other data to that time of validity. The inertial solution update rate should be constrained to an integer multiple of the GPS measurement update rate. Due to data-link capacity constraints, the data fusion processor will output most data at a slower rate than it inputs it.

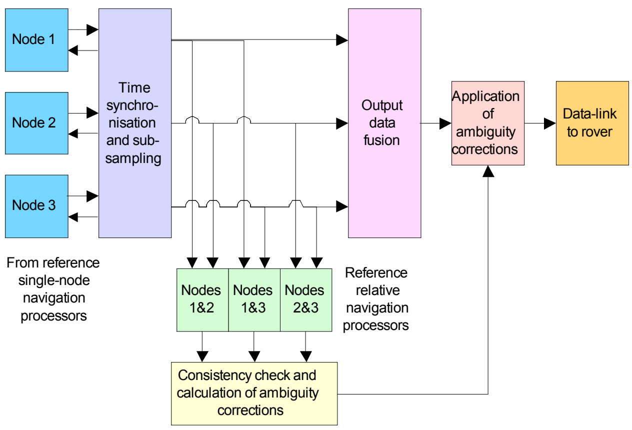

Figure 3: Simplified flowchart for reference data fusion processor (excludes integrity monitoring)

The output data fusion function combines the corrected inertial navigation solution, GPS pseudo-range and ADR measurement data, Kalman filter float ambiguity, range bias and ionosphere estimates, the full error covariance matrix and the INS corrections from each of the nodes to simulate the output of a virtual INS/GPS navigation system located at the centroid of the *M* nodes.

Equal weighting of data from the various INS/GPS nodes in the data fusion algorithm is preferred because an equally weighted navigation solution, measurement residual set and Kalman filter is implicitly describing the centroid body frame without the need to apply lever arm corrections, which are subject to measurement error. The main disadvantage is that equal weighting does not generally minimise the noise. However, weighting the GPS measurement data, navigation solutions and Kalman filter state estimates according to their variances/covariances introduces the problem that the Kalman filter and navigation solution weighting does not match the history of the GPS measurement weighting. This could impact the validity of the rover-reference relative navigation filter's error covariance matrix, which is highly undesirable from an integrity perspective. Thus, equal weighting is the lowest risk option.

The reference relative navigation processors each calculate an ambiguity-fixed relative navigation solution between two of the reference system nodes. Calculating relative navigation solutions for all pairs of nodes rather than just the independent pairs provides for consistency checks to be performed to verify the integrity of those solutions. This partitioned approach is more processor-efficient than a centralised relative navigation processor which calculates the relative positions of all the nodes simultaneously, particularly for ambiguity fixing, as the centralised processor load scales non-linearly with the number of nodes.

Each relative navigation processor comprises an EKF estimating double-differenced float ambiguities, together with the relative position and velocity errors, and an ambiguity fixing algorithm, such as the least-squares ambiguity decorrelation adjustment (LAMDA) method [19]. Range biases and ionosphere propagation delays are assumed to cancel between INS/GPS nodes on the same host vehicle.

Following the application of consistency checks to the reference relative navigation processors' ambiguity estimates, ambiguity corrections are calculated and then applied to both the ADR measurement data and the float ambiguity estimates from the output data fusion function prior to transmission over the data link to the rover vehicle.

5. RELATIVE NAVIGATION

The key design decision for the rover-reference relative navigation processor is whether to implement a partitioned or a differenced architecture within the extended Kalman filter that forms the processor's core. In the partitioned architecture, separate measurements are input and separate states estimated for the rover and reference vehicle. The correlation between rover and reference states is modelled in the off-diagonal elements of the error covariance matrix, **P**. In the differenced architecture, all measurements and states are differenced between rover and reference. Note that, in both cases, the ambiguities are differenced between rover and reference before input to fixing algorithm as the satellite phase biases must be cancelled for the integer constraint to be applicable.

Figures 4 and 5 depict simplified flowcharts for the rover-reference relative navigation processor using the partitioned and differenced architecture, respectively. Integrity monitoring, IODE/IODC consistency checks and Kalman filter tuning and initialisation data are not shown.

The key advantage of the differenced architecture over the partitioned is that half the number of Kalman filter measurements and states are used, which results in a processing load about 8 times lower [11].

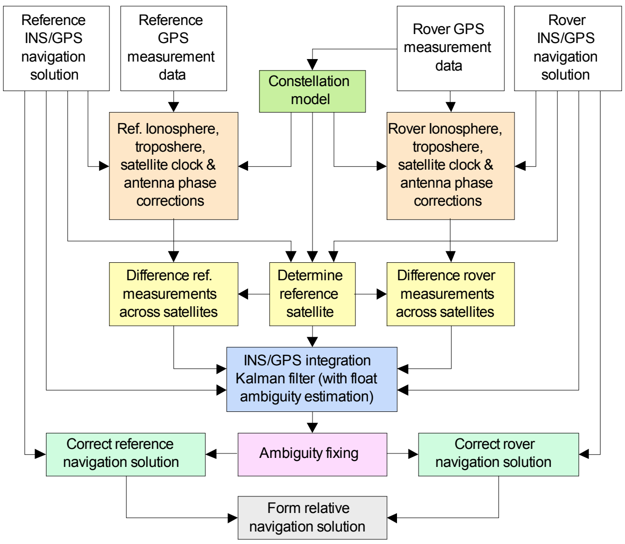

Figure 4: Simplified flowchart for rover-reference navigation processor with partitioned architecture (excludes integrity monitoring)

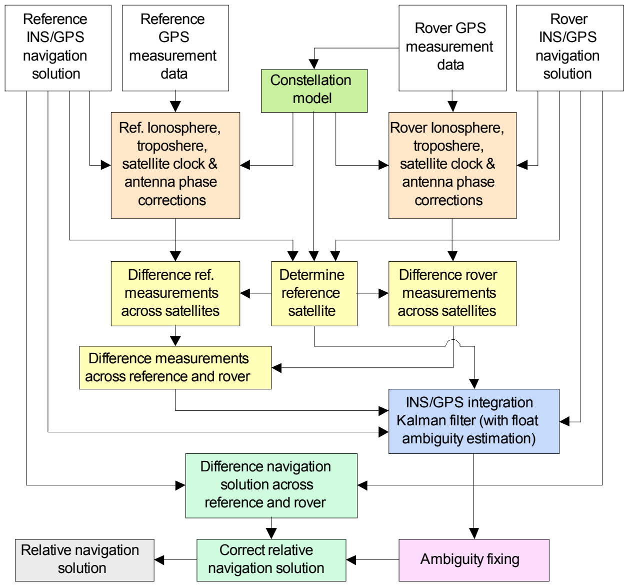

Figure 5: Simplified flowchart for rover-reference navigation processor with differenced architecture (excludes integrity monitoring)

However, the greater flexibility of the partitioned architecture also brings benefits:

- Rover measurements can still be processed where there is a data-link interruption.
- It can handle the tracking of different satellites by the rover and reference GPS user equipment, for example when a rover manoeuvre results in airframe blockage of a line of sight or where a low elevation satellite is above the masking angle for one host vehicle but not the other.
- Where the integrity monitoring system identifies a fault, it is much easier to determine whether it has arisen in the rover or reference data.
- Rover and reference GPS measurement data does not have to be time-synchronised; instead, separate rover and reference measurement updates may be performed with the system propagation phase of the EKF bridging the gap.
- Rover and reference GPS measurement data may be processed at different rates; this can be useful where the rover GPS user equipment outputs at a faster rate than the data link can provide reference GPS measurements.

The discussion on processing L1 and L2 measurement data separately or combined for the single-node navigation processors in Section 3 also applies to the relative navigation processor.

However, the choice of which states to estimate is different. Essentially, the relative navigation processor exploits correlations in the errors exhibited by the rover and reference navigation systems to generate a more accurate relative navigation solution than would be obtained simply by differencing the two independent navigation solutions. The correlated errors are the GPS ephemeris, satellite clock, ionosphere and troposphere errors. The receiver clock and IMU errors within the reference and rover are independent, so the relative navigation EKF will not improve upon the calibration of these errors within the single-node EKFs. Consequently, the relative navigation EKF states should be limited to position and velocity error, range biases, ionosphere propagation delays and float ambiguities. GPS states and measurements must be differenced across satellites in order to cancel out the receiver clock errors.

The position and velocity error states estimate the errors in the corrected rover and fused reference navigation solutions. As in any federated integration architecture [11, 15] any corrections fed-back to the inertial navigation equations function from one filter must also be applied in all the other filters that estimate the INS errors. This is because corrections change the true value of the states estimated by the Kalman filters. Here, the feedback is generated within the single-node navigation processors, so the corrections must be transmitted from the nodes and also applied within the relative navigation processor.

6. INTEGRITY

An integrity requirement for a navigation solution is defined in terms of a maximum probability over a certain time interval that the position error exceeds a certain threshold without the user being alerted. For demanding applications, integrity monitoring is required to ensure that faults resulting in the position error exceeding the threshold are detected. Furthermore, on detection of a fault, a fault-free navigation solution should be provided using the remaining data, where possible; a concept known as fault isolation.

An inherent feature of INS/GPS integration is the use of a Kalman filter-based estimation algorithm to determine INS and IMU corrections from the GPS measurements [11]. It is also proposed here to use the Kalman filters to determine the GPS carrier-phase float ambiguities, noting that other methods are available and that a fixing algorithm, such as LAMBDA [19] must also be used to apply the integer constraint.

However, using a Kalman filter makes it more difficult to isolate faulty measurement streams from the navigation solution than where a snapshot GPS position solution is used. This is because, unless the fault has been detected immediately, the navigation solution will have already have been contaminated by faulty data prior to the detection of the fault.

Fault isolation can be achieved by performing a complete reset of the inertial navigation solution(s), Kalman filter(s) and ambiguity resolution. However, there will then be a delay before the filtered navigation solution meets the integrity and accuracy requirements. This is likely to breach the continuity requirement.

Therefore, for an INS/GPS system to meet demanding accuracy, integrity, continuity and availability requirements (see [11] for definitions), it is necessary to maintain parallel navigation solutions for the system to revert to in the event of a fault [20, 21]. To protect against faulty GPS signals, a series of parallel navigation solutions must be maintained, each under the hypothesis that data from one of the satellites must not be used. These are sometimes known as H1 hypotheses, with the fault-free hypothesis known as H0. When a fault is detected in one of the satellite signals, the output is switched from the H0 hypothesis navigation solution to the appropriate H1 hypothesis solution. Figure 6 illustrates this for a single-node navigation processor. Parallel reference data fusion and relative navigation processors must also be maintained for each faulty satellite H1 hypothesis.

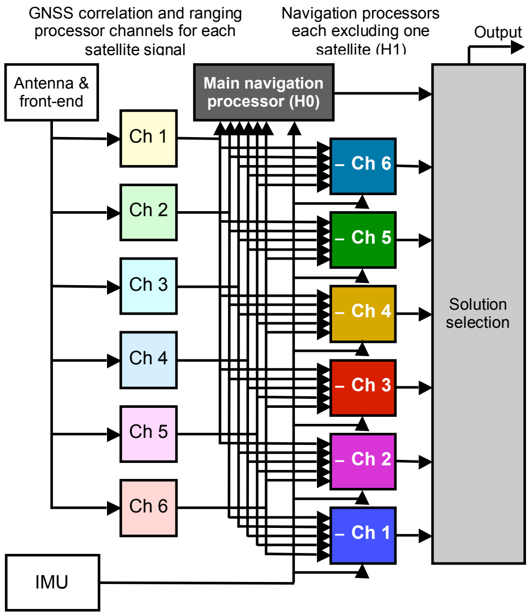

Figure 6: Multiple-hypothesis parallel INS/GNSS navigation processors for robustness against GNSS faults (after [11])

Where multiple reference nodes are implemented, the integrity monitoring system should also protect against faults in a single-node GPS user equipment, IMU or processor. Therefore, H1 hypothesis parallel solutions should also be maintained under the hypothesis that data from one of the reference nodes must not be used. These require replication of only the output data fusion, calculation of ambiguity corrections and application of ambiguity corrections functions of the reference fusion processor (see Figure 3), together with the complete rover-reference relative navigation processor. Note that data for all of the H1 hypotheses must be transmitted on the reference to rover data-link in addition to the H0 data.

The remaining hypotheses are collectively known as H2. This covers faults occurring on the rover, the data-link and multiple cases of satellite and/or reference node fault. When H2 faults occur, the relative navigation solution must be declared unavailable. Note that if multiple INS/GPS nodes are also implemented on the rover, a H1 hypothesis should be maintained for each rover node fault as well.

In the event that the integrity monitoring system identifies the H0 solution as contaminated by faulty data, it is not sufficient to simply switch to the appropriate H1 hypothesis solution. It is also necessary to ensure that the H1 solution is fault-free. This concept is known as fault exclusion. Therefore, the integrity monitoring architecture must be replicated across all parallel navigation processors.

Moving on from fault exclusion to fault detection, QinetiQ has identified 48 different integrity monitors which can potentially contribute to the overall performance of a relative carrier-phase INS/GPS system of the type discussed in this paper. They are listed below in order of location within the processing chain. These monitors are currently being assessed, with opportunities for improvement in each design identified. Note that a number of individual monitors used for LAAS are grouped into three broad categories [14]: data quality monitor (DQM), measurement quality monitor (MQM) and signal quality monitor (SQM). These are marked in brackets on the names of the relevant monitors below.

The single-node navigation processor integrity monitors are:

- IMU output range checks [11];
- INS solution range checks [11];
- Abnormal dynamics test;
- GPS signal to noise level;
- GPS ephemeris check (DQM) [14];
- General GPS data (DQM) [14];
- GPS carrier phase lock (MQM) [14];
- GPS carrier frequency lock;
- GPS carrier parity (MQM) [14];
- GPS evil waveform detection (SQM) [14];
- GPS ionosphere anomaly monitoring;
- GPS code-carrier divergence (SQM) [14];
- GPS carrier acceleration and step test (MQM) [14];
- GPS cycle slip detector [22];
- Receiver autonomous integrity monitoring (RAIM) [23]:
    - Single-epoch code;
    - Single-epoch carrier change (delta-range); sometimes known as relative RAIM [24];
    - Code over time;
    - Carrier over time;
- GPS code acceleration and step test;
- Kalman filter measurement innovation filtering [11];
- Kalman filter measurement innovation sequence monitoring [11];
- Kalman filter state range checks [11];
- Equipment failure monitor.

The reference data fusion processor integrity monitors are:

- Reference node ephemeris consistency check;
- GPS message parity check;
- Reference node solution consistency check;
- Reference node measurement consistency check;
- Reference node IODE and IODC consistency check;
- Reference node signal to noise consistency check;
- Message field range test (MFRT) [14];
- Multiple reference consistency check (MRCC) [14];
- Sigma mean monitor (SMM) [14];
- Reference node relative ambiguity consistency check;
- Reference node relative position/velocity consistency check;
- Equipment failure monitor;
- Data-link monitor.

The rover-reference relative navigation processor integrity monitors are:

- Rover-reference ephemeris consistency check;
- Data-link cyclic redundancy check (CRC);
- Rover-reference range check;
- Reference data latency check;
- Rover-reference ionosphere estimate consistency check;
- Rover-reference range bias estimate consistency check;
- Kalman filter measurement innovation filtering [11];
- Kalman filter measurement innovation sequence monitoring [11];
- Kalman filter state range checks [11];
- Ambiguity fix residual monitor;
- Probability of false ambiguity fix monitor;
- Equipment failure monitor.

QinetiQ has conducted a comprehensive failure modes and effects analysis for relative carrier-phase INS/GPS. Unlike previously published studies [25, 26], the causes of failures and the effects they have on the various signals input to the navigation processors have been separately tabulated and the connections between them fully determined. For example, the causes include satellite hardware faults, multipath and reference node IMU failure, while the effects include GPS carrier ramp error, unavailable GPS navigation data message and IMU noise burst. To date, 32 consolidated failure causes and 46 failure effects have been identified.

Each integrity monitor is being matched to one or more failure mode effects in the overall relative INS/GPS system. This will be followed by an audit of monitors against failure effects to remove any duplication and identify any gaps in the monitoring.

The final real-time feature of an integrity monitoring architecture for safety critical applications is the requirements monitor. This determines whether the relative navigation solution currently meets the accuracy, integrity and continuity requirements. This is done by ensuring that a sufficiently high proportion of the position error distribution lies within the relevant alert limits. Unfortunately, the true error distribution is non-Gaussian while navigation processors based on Kalman filters model error distribution as Gaussian. Therefore, it is essential to ensure that the Gaussian approximation to the error distribution is a conservative overbound of the true error. This means that the proportion of the Gaussian approximation lying outside the limits must always be greater than or equal to the proportion of the true distribution outside those limits.

Various overbounding methods have been devised since the initial approach by deCleene was published [27]. A drawback of most methods is that they cannot be applied to overbound time-correlated errors. This disadvantage has been noted in [28], where a spherically symmetric overbound was introduced. The bound was originally formulated in terms of impulse response coefficients. Recently, QinetiQ has reformulated the bound in a state-space framework [29], making it easily applicable to linear systems such as Kalman filters.

7. FURTHER WORK

Work is currently ongoing to identify improvements to the integrity monitors. To date, innovative approaches have been identified for improving the performance of the following monitors:

- IMU output range checks;
- INS solution range checks;
- GPS signal to noise level;
- GPS ionosphere anomaly monitoring;
- GPS code-carrier divergence;
- GPS cycle slip detector;
- RAIM;
- Equipment failure monitor;
- Reference node solution consistency check;
- Reference node signal to noise consistency check.

In combining ambiguity-fixed carrier-phase GPS with high integrity requirements, it is necessary to determine the impact of the fixing process on the solution integrity. This is because applying an ambiguity fix changes the position error distribution from single-modal to multi-modal. If the probability of correct fix (PCF) is not sufficiently high or cannot be determined with sufficient precision, the integrity performance of the fixed solution can be worse than that using the float ambiguities (i.e. where the integer constraint is not applied).

A partial solution to this problem is the recent proposal to include the probability of an almost-correct fix (PAF) in the integrity calculation [30]. However, the problem of developing a robust statistical basis for calculation of PCF (and PAF) still remains. QinetiQ is currently investigating this alongside an innovative approach to the fix/ don't fix decision.

For a navigation system to be certified for use in safety critical applications, it must be verified that it meets the requirements. This is done by determining whether the accuracy, integrity and continuity requirements are met at all locations where the system will be used and for all GPS constellation configurations. The proportion of times and locations where the requirements are met is known as the availability. This availability modelling must account for all failures modes and effects, weighted by their probability of occurrence.

For a GPS single-point (or "snapshot") navigation solution, the accuracy, integrity and continuity tests can be performed analytically. So availability modelling may proceed simply by repeating these tests over a spatial and temporal grid as QinetiQ has demonstrated previously [4]. However, this approach does not extend to GPS and INS/GPS filtered navigation solutions. Instead, a Monte-Carlo approach is usually required. Performing an integrity test at a single point in time and space typically requires 100 times as many runs as the expected number of occurrences of the integrity failure. For a navigation system with an integrity requirement for the missed fault detection probability to be within 10^{-7} over a defined period, a pure Monte-Carlo approach would necessitate around 10^{9} runs per point in a spatial and temporal grid. This is clearly computationally infeasible.

QinetiQ is instead developing a hybrid simulation and analysis approach to aid with assessing the performance of high integrity navigation systems. For a given system configuration consisting of satellite constellation and user equipment location, a technique is being sought to overbound the true navigation system error distribution so that questions concerning accuracy and integrity can be addressed without recourse to "brute-force" Monte-Carlo simulations. A key to the approach is the consolidation of failure modes and effects discussed in Section 6.

As the model is to be a conservative overbound of the truth, it should never pass an accuracy, integrity or continuity test when the true navigation system would not. By repeating these tests over a spatial and temporal grid, a lower bound on the availability performance may be obtained. This is sufficient for certifying the system as it is only necessary to show that the availability exceeds the requirement; there is no need to determine what the availability actually is.

8. CONCLUSIONS

The preferred processing architecture for high-integrity carrier-phase relative INS/GPS distributes the processing between single INS/GPS nodes aboard rover and reference, a reference data fusion algorithm and rover-reference relative navigation processor.

Separate processing of L1 and L2 pseudo-range and ADR measurements in both single-node and relative Kalman filters is recommended. Calibration of receiver clock errors, INS attitude errors and accelerometer and gyro biases should be in the single-node filters only. Position and velocity errors, GPS range biases, ionosphere propagation delays and float ambiguities should be estimated in both single-node and relative filters, with the relative filter initialised using information from the single-node filters.

Where multiple INS/GPS nodes are implemented on the reference vehicle for added integrity, it is recommended that their outputs be fused prior to transmission to the rover. This brings the benefits of reduced data transmission lags, fewer carrier-phase ambiguities to resolve aboard the rover and no need for the rover's navigation system to model flexing lever arms between reference nodes. However, a simple averaging of reference node data leads to non-integer-wavelength carrier-phase ambiguities. QinetiQ's patent-pending solution is to fix the relative ambiguities of the reference node and use this to correct the ambiguities of the fused reference ADR measurements.

To achieve fault exclusion when a fault is not immediately detected by the integrity monitoring system, parallel navigation processors must be maintained, each under the hypothesis that data from one of the GPS satellites or reference INS/GPS nodes should not be used. Should a fault in one of the satellite signals or reference nodes be discovered after it has impacted on the main navigation solution, the system should revert to a reversionary solution.

QinetiQ has conducted a comprehensive failure modes and effects study for carrier-phase relative INS/GPS, mapping 32 consolidated failure causes to 46 failure effects. To detect these failures, 48 types of integrity monitor have been identified and improvements proposed for some of them. QinetiQ is currently investigating improving the integrity of ambiguity fixing and developing a hybrid analytical and simulation-based performance model for determining a lower bound to relative INS/GPS solution availability.

REFERENCES

1. S. Khanafseh, B. Pervan and G. Colby, *Carrier Phase DGPS for Autonomous Airborne Refueling,* ION GNSS 2005

2. J. Goddard, S. Mole and D. Nash, "Application of Shipboard Relative GPS to Automatic Recovery of STOVL Aircraft," *Proc. NAV05,* London, November 2005.

3. C. Mather, S. Mole, A. Macaulay and J. Goddard, "Performance of Integrity Monitoring Techniques for Shipboard Relative GPS Landing Systems," *Proc. ION GNSS*, September 2005.

4. C. Offer, P. D. Groves, A. Macaulay, D. Nash and C. Mather, "Use of Inertial Integration to Enhance Availability for Shipboard Relative GPS (SRGPS) Landing Systems," *Proc. ION GNSS*, September 2006.

5. P. D. Groves, C. Mather and A. Macaulay, "Development of non-coherent deep INS/GPS integration for optimal signal to noise performance", *Proc. ION GNSS*, September 2007.

6. C. Mather, P. D. Groves and M. R. Carter, "A man motion navigation scheme using high sensitivity GPS, MEMS IMU and auxiliary sensors", *Proc. ION GNSS*, September 2006.

7. P. D. Groves, G.W. Pulford, A. Littlefield, D. Nash and C. Mather, "Inertial navigation versus pedestrian dead reckoning: Optimizing the integration", *Proc. ION GNSS*, September 2007.

8. P. D. Groves and D. C. Long, "Combating GNSS Interference with Advanced Inertial Integration," *Journal of Navigation*, Vol. 58, No. 3, 2005, pp. 419–432.

9. P. D. Groves, R. J. Handley and A. R. Runnalls, "Optimising the Integration of Terrain Referenced Navigation with INS and GPS," *Journal of Navigation*, Vol. 59, No. 1, 2006, pp. 71–89 and *Proc. ION GNSS*, September 2004.

10. P. D. Groves, "Optimising the Transfer Alignment of Weapon INS," *Journal of Navigation*, Vol. 56, No. 3, 2003, pp. 323–335.

11. P. D. Groves, *Principles of GNSS, Inertial and Multi-sensor Integrated Navigation Systems*, Artech House, January 2008.

12. E. D. Kaplan and C. J. Hegarty (eds), *Understanding GPS: Principles and Applications*, 2nd edition, Artech House, 2006.

13. B. W. Parkinson and J. J. Spilker, Jr., *Global Positioning System: Theory and Applications*, Vols. I and II, AIAA, 2006.

14. G. Xie, S. Pullen, M. Luo, P. L. Normark, D. Akos and J. Lee, "Integrity Design and Updated Test Results for the Stanford LAAS Integrity Monitor Testbed", *Proc. ION 57th Annual Meeting*, June 2001

15. N. A. Carlson, "Federated Filter for Distributed Navigation and Tracking Applications," *Proc. ION 58th AM*, Albuquerque, NM, June 2002, pp. 340–353.

16. D. H. Titterton, and J. L. Weston, *Strapdown inertial navigation technology*, 2nd Edition, IEE, 2004.

17. M. S. Grewal, L. R. Weill and A. P. Andrews, *Global Positioning Systems, Inertial Navigation, and Integration*, 2nd edition, Wiley, 2007.

18. C. R. Offer, *Phase Based Measurement Corrections*, Patent application WO 2008/075002 A1, published 26/6/2008, priority date 19/12/2006.

19. P. J. G. Teunissen, P. J. De Jonge and C. C. J. M. Tiberius, "Performance of the LAMBDA Method for Fast GPS Ambiguity Resolution," *Navigation: JION*, Vol. 44, No. 3, 1997, pp. 373–383.

20. J. Diesel and S. Luu, "GPS/IRS AIME: Calculation of Thresholds and Protection Radius Using Chi-Square Methods," *Proc. ION GPS*, September 1995.

21. Brenner, M., "Integrated GPS/Inertial Fault Detection Availability," *Navigation: JION*, Vol. 43, No. 2, 1996, pp. 111–130.

22. O. L. Colombo, U. V. Bhapkar and A. G. Evans, "Inertial-Aided Cycle-Slip Detection/ Correction for Precise, Long-Baseline Kinematic GPS," *Proc. ION GPS*, September 1999.

23. R. G. Brown and P. W. McBurney, "Self-Contained GPS Integrity Check Using Maximum Solution Separation as the Test Statistic," *Proc. ION GPS-87*, Colorado Springs, CO, September 1987, pp. 263–268.

24. T. Walter, J. Blanch, P. Enge, B.s Pervan and L. Gratton, "Future Architectures to Provide Aviation Integrity", *Proc. ION NTM*, January 2008.

25. US Dept. of Transportation, Federal Aviation Administration, *Category I Local Area Augmentation System Ground Facility Specification*, FAA-E-2937A, 17 April, 2002.

26. U. I. Bhatti and W. Y. Ochieng, "Failure Modes and Models for Integrated GPS/INS Systems," *Journal of Navigation*, Vol. 60, No. 2, 2007, pp. 327–348.

27. B. DeCleene, "Defining Pseudorange Integrity-Overbounding," *Proc. ION GPS 2000*, Salt Lake City, UT, September 2000, pp. 1916–1924.

28. J. Rife and D. Gebre-Egziabher, "Symmetric Overbounding of Correlated Errors," *Navigation: JION*, Vol. 54, No. 2, 2007, pp. 109-124.

29. G. W. Pulford, "A Proof of the Spherically Symmetric Overbounding Theorem For Linear Systems," *Navigation: JION*, 2008, to appear.

30. S. Wu, S. R. Peck. R. M. Fries and G. A. McGraw, "Geometry Extra-Redundant Almost Fixed Solutions: A High Integrity Approach for Carrier Phase Ambiguity Resolution for High Accuracy Relative Navigation," *Proc. IEEE/ION PLANS*, May 2008.

This paper was first presented at ION GNSS 2008, 16-19 September 2008, Savannah, GA, USA